Hybrid Forecasting for Airline Revenue Management in Semi-Restricted Fare Structures

By

Michael H. Reyes

B.S., Civil Engineering
The University of Texas at Austin, 2004

SUBMITTED TO THE DEPARTMENT OF CIVIL AND ENVIRONMENTAL ENGINEERING IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE IN TRANSPORTATION
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

JUNE 2006

Signature of Author: ……………………………………………………………………………………………….
Department of Civil and Environmental Engineering
May 9, 2006

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

Certified by: ……………………………………………………………………………………………………….
Joseph M. Sussman
JR East Professor of Civil and Environmental Engineering
Thesis Reader

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

Hybrid Forecasting for Airline Revenue Management in Semi-Restricted Fare Structures

by

Michael H. Reyes

Submitted to the Department of Civil and Environmental Engineering on May 9, 2006 in Partial Fulfillment of the Requirements for the Degree of Master of Science in Transportation

ABSTRACT

In recent years, the airline industry has seen diminished performance of traditional Revenue Management (RM) systems largely due to the growth of Low Cost Carriers and the increased use of their "simplified" fare structures. With the removal of many of the fare restrictions essential to RM systems, standard demand forecasters can no longer segment demand and passengers are able to book air travel in fare classes lower than their actual willingness to pay. These "semi-restricted" fare structures typically contain several homogenous fare classes undifferentiated except by price, as well as several distinct fare classes with unique combinations of booking restrictions and advance purchase requirements.

This thesis describes "Hybrid Forecasting" (HF) – a new technique which separately forecasts "product-oriented" demand using traditional forecasting methods, and "price-oriented" demand for passengers who purchase only in the lowest priced fare class available when booking. The goal of this thesis is two-fold: to first measure the potential benefit of Hybrid Forecasting in terms of network revenue in semi-restricted fare structures, and then to measure potential improvements to Hybrid Forecasting. "Path Categorization" attempts to improve revenues by exploiting the expected higher level of passenger willingness-to-pay for non-stop service versus connecting service. And "Fare Adjustment" accounts for passenger sell-up behavior from lower to higher fare classes, and is applied within an RM system's seat inventory optimizer.

Experiments with the Passenger Origin-Destination Simulator demonstrate that HF in these semi-restricted fare structures can improve an airline's network revenue by approximately 3% compared to traditional forecasting methods. This improvement grows by 0.25% with Path Categorization, by 1% with Fare Adjustment, and by up to 2.5% over Hybrid Forecasting alone with Path Categorization and Fare Adjustment together – all significant impacts on an airline's network revenue. Though these results are encouraging, the revenue gains of these new RM forecasting methods are still not enough to offset the revenue loss associated with the easing of traditional fare class restrictions.

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

Thesis Reader: Dr. Joseph M. Sussman
Title: JR East Professor of Civil and Environmental Engineering

Acknowledgements

First, I must thank Dr. Peter Belobaba for supervising both my research and this thesis. He has heightened my understanding both of Revenue Management and the airline industry, in general, and I am enormously grateful for his agreeing to be my supervisor while at MIT.

I am also indebted to Dr. Joseph Sussman both for agreeing to read this thesis, as well as supervising me throughout my MIT career. I must thank him for allowing me to serve as Teaching Assistant for his Introduction to Transportation Systems course.

Graduate study at MIT can be an expensive proposition, and I am honored that the Federal Highway Administration's Universities and Grants Program selected me for an Eisenhower Graduate Fellowship in 2004. Without the support of Dr. Ilene Payne, Gwen Sutton, Henry Murdaugh, Camron Ranje, Walter Woods, and everybody else at the FHWA in the past two years, I would never have been able to attend the Institute.

Furthermore, I must thank the airlines in the Passenger Origin-Destination Simulator (PODS) Consortium for making this research possible through their financial support, as well as Craig Hopperstad for explaining to me the workings of the simulator as I repeatedly attempted to break it.

I am grateful to my fellow PODS students in the MIT International Center for Air Transportation for helping me learn about the simulator, including Maital Dar, Val Soo, Thierry Vanhaverbeke, and Greg Zerbib. Of course, I must also acknowledge Stephane Bratu – a former PODS student and MIT alum – for helping me understand RM. And I also thank my friends at Sabre Airline Solutions who first introduced me to the science of Revenue Management.

In the MST program, I thank Lucile, Drew, Chris, Edgar, Owen, Lev, Travis, and everybody else for their encouragement the last two years.

Without my professors at the University of Texas I never would have been prepared for the rigors of graduate school. Special thanks go to Dr. Kara Kockelman, Chris Frazier, and Dr. Jorge Prozzi for allowing me to work for them as an undergraduate.

And finally, I must thank my support network from Texas: Jill, Mom, Dad, Joseph, and of course, my Grandma. If not for them, I never would have made it into or out of MIT.

Author's Biography

Born in Austin – a beautiful city in the Texas Hill Country – the author was raised 90 minutes to the south in San Antonio. There he attended Tom C. Clark High School before returning to Austin in 2000 to study Civil Engineering at the University of Texas. While an undergraduate, he worked for the City of Boerne, Texas, for Pape-Dawson Engineers, Inc., and as a student researcher at the University. His areas of research included pavement design with Dr. Jorge Prozzi, and urban land use patterns with Dr. Kara Kockelman and Chris Frazier. In May of 2004, he graduated from the University with Highest Honors.

In September of that year, he entered the Massachusetts Institute of Technology pursuing a Master's Degree in Transportation on an Eisenhower Graduate Fellowship from the Federal Highway Administration. Shortly thereafter, he specialized in air transportation studies, and took the following courses:

- Introduction to Transportation Systems with Dr. Joseph Sussman
- Transport Demand and Economics with Dr. Moshe Ben-Akiva
- Transportation Policy, Strategy, and Management with Dr. Joseph Coughlin
- Carrier Systems with Dr. Cynthia Barnhart and Dr. Nigel Wilson
- Transportation Flow Systems with Dr. Amedeo Odoni and Dr. Patrick Jaillet
- Planning and Design of Airport Systems with Dr. Amedeo Odoni and Dr. Richard de Neufville
- The Global Airline Industry with Dr. Peter Belobaba
- Airline Management with Dr. Peter Belobaba
- Airline Schedule Planning with Dr. Cynthia Barnhart
- Demand Modeling with Dr. Moshe Ben-Akiva
- Competitive Decision Making and Negotiation with Dr. Gordon Kaufman

In the Fall of 2005, he served as Teaching Assistant for Dr. Joseph Sussman's Introduction to Transportation Systems course. Also in 2005, he began working as a Research Assistant with MIT's Passenger Origin-Destination Simulator (PODS) Consortium led by Dr. Peter Belobaba. For the PODS Consortium, he has presented research in Paris, France (February, 2006), and his research interest is demand forecasting for airlines in less-restricted fare structures. He anticipates ending his Master's career at MIT by presenting at the PODS meeting in Santiago, Chile (May, 2006), as well as presenting at the Airline Group for the International Federation of Operational Research Societies (AGIFORS) Revenue Management Study Group meeting in Cancun, Mexico (May, 2006).

Table of Contents

List of Figures

List of Tables

List of Abbreviations

AP	Advance Purchase
ASM	Available Seat Mile
AT	Adaptive Threshold
BC*i*	Base Case *i*
BP	Bid Price
CAB	Civil Aviation Board
DAVN	Displacement Adjusted Virtual Nesting
DFW	Dallas-Fort Worth International Airport
DM	Demand Multiplier
DP	Dynamic Programming
DWM	Decision Window Model
EMSR	Expected Marginal Seat Revenue
FA	Fare Adjustment
FCFS	First Come First Served
FCYM	Fare Class Yield Management
FRAT	Fare Ratio
FRAT5	Fare Ratio at which 50% of passengers demonstrate sell-up
FT	Fixed Threshold
GVN	Greedy Virtual Nesting
HF	Hybrid Forecasting
HF1	Method of classifying historical bookings using the “path” rule
HF2	Method of classifying historical bookings using the “airline” rule
HF3	Method of classifying historical bookings using the “market” rule
HBP	Heuristic Bid Price
IAP0	Classifying historical AP bookings as price-oriented
IAP1	Classifying historical AP bookings as product-oriented
IAP2	Conditionally classifying historical AP bookings as product-oriented
KI	Karl Isler Fare Adjustment
LCC	Low Cost Carrier
MR	Marginal Revenue Fare Adjustment
MSP	Minneapolis-St. Paul International Airport
NetBP	Network Bid Price
OD	Origin-Destination, or Origin to Destination
PODS	Passenger Origin-Destination Simulator
PCAT1	Path where carrier provides superior path quality to competition
PCAT2	Path where carrier provides same path quality as competition
PCAT3	Path where competition provides superior path quality to carrier
PE	Price Elasticity
ProBP	Prorated Bid Price
R*i*	Booking Restriction *i*, where *i* = 1, 2, or 3
RM	Revenue Management
SAS	Scandinavian Airlines
WTP	Willingness-to-Pay
YM	Yield Management

1 Introduction

In the nascent years of airline Revenue Management (**RM**) systems, American Airlines once simplistically described the developing practice as "selling the right seats to the right customers at the right prices"[1]. This was, and still is, the goal of Revenue Management, though our understanding of the problem and our approaches to solving it have evolved tremendously in a relatively short amount of time.

A generation ago, RM could have been considered a narrow area of interest to academics and airline operations enthusiasts; it was somewhat of a curiosity in the heavily regulated industry where airlines had minimal control over fares and booking methods. Today, RM is an indispensable tool, as nearly every carrier in the world seeks to maximize passenger revenue (and thus profits) by extracting fares at customers' highest willingness-to-pay (**WTP**).

Following deregulation of the US airline industry in 1978, airlines faced two choices: either adaptation to a new business environment – one without artificial limits on competition – or obsolescence. And just as the nimble airlines once developed creative new RM approaches to confront wholesale changes in the business of providing air transportation to the public, today's carriers are challenged to adapt to a new competitive environment – one where the assumptions previously made about customers' booking habits have been invalidated.

The goal of this thesis is to explore a new method of demand forecasting for airline RM in a changing competitive environment: hybrid forecasting (**HF**) in semi-restricted fare structures. Specifically, we will present and evaluate HF as a new approach to predicting passenger demand and maximizing passenger revenue when the top (and sometimes only) priority for many consumers is the minimization of spending on airfare.

In this thesis, we will describe HF, and then demonstrate the potential revenue gains of HF alone, and in conjunction with selected other RM strategies including "Fare Adjustment" (**FA**) and "Path Categorization" (**PCAT**). We will employ a simulation approach, utilizing the Passenger Origin-Destination Simulator (**PODS**) originally developed at the Boeing Company in order to model the airline booking process with competing carriers seeking to maximize passenger revenues over a given network.

1.1 Deregulation and Revenue Management Background

The first commonly accepted application of airline RM – then known as yield management (**YM**) – occurred in the early 1970's when BOAC (later known as British Airways) introduced two fare products for a single inventory of seats.[2] However, it is worth noting that operations researchers had been musing on overbooking – the practice of deliberately selling more seats than physically available in order to mitigate financial damage by absentee passengers – since the 1960's; and airlines

[1] Smith, B. C., J. F. Leimkuhler, R. M. Darrow. 1992. Yield management at American Airlines. *Interfaces.* Volume 22, Issue 1, pp. 8-31.
[2] McGill, J. I., G. J. van Ryzin. 1999. Revenue management: research overview and prospects. *Transportation Science.* Volume 33, Issue 2, pp. 233–256.

may have been furtively engaging in the practice before then, though the potential for public relations fallout at the time likely precludes our latter day knowledge of such strategies.[3] Regardless of these early contributions to the RM field, it is widely acknowledged that deregulation of the industry was the catalyst for tremendous gains in the field of RM.

In the United States before deregulation, the Civil Aviation Board (CAB) set airline fares for the entire industry, denying individual airlines to right to increase (or decrease) fare levels except in the case where financial losses could be demonstrated. Concerned that the CAB's control of fares (as well as routes) inhibited growth of the nation's air transportation network and permitted competitive inefficiencies, in 1978 Congress deregulated the industry and granted carriers control over their respective product offerings.[4] Deregulation allowed not only existing carriers freedom to price flights in creative new ways, but relaxed barriers to entry and allowed new competitors to begin flying – competitors with much lower cost structures, which allowed them to charge lower fares to passengers.

After deregulation, pricing structures quickly evolved to reflect demand for specific origin-destination (**OD**) market pairs, instead of the per-mile pricing previously enforced by the CAB[5]. As prices dropped in many markets to reflect increased competition for passengers, many airlines – exposed to the open market for the first time – learned that profit maximization would theoretically occur when each individual consumer paid exactly the maximum amount that he or she was willing to pay for airfare. Thus, the idea of differential pricing moved from operations research theory to heavy application within major airlines.

The basis for RM over the last twenty years has been the distinction between two general types of airline passengers and forecasting the demand of each for air transportation: leisure and business travelers. In general, leisure passengers tend to know travel plans weeks in advance of departure, are flexible in terms of departure and arrival times when booking, and have a low WTP for airfare. In contrast, business passengers often travel with very little notice, exhibit schedule inflexibility, and are willing to pay for the convenience of last minute travel.

In terms of revenue paid per passenger per mile, or yield, fares typically paid by leisure passengers are lower yield than those in the business segment. And though the obvious revenue maximization strategy for any airline may seem to be yield maximization (or filling a plane entirely with high yield business passengers), in practice the marriage of demand and supply is an awkward one. To an airline, demand is represented by an ever changing number of passengers, each with a different expectation of the cost of airfare; while supply is a fixed inventory of aircraft seats, unchangeable in the short term.

The general supply-demand relationship for this situation – on the simplest basis of a single flight leg with no connecting traffic considered - is shown in Figure 1. Note that this Figure assumes not enough high yield business demand exists to fill the aircraft to capacity, but that the overall number of number of people willing to pay

[3] Rothstein, M. 1985. O.R. and the airline overbooking problem. *Operations Research.* Volume 33, Issue 2, pp. 237–248.
[4] General Accounting Office. 1999. Airline deregulation: changes in airfares, service quality, and barriers to entry. Report to Congressional Requesters. GAO/RCED-99-92. Washington, D.C.
[5] Pickrell, D. 1991. The regulation and deregulation of US airlines. *Airline deregulation: international experiences.* Ed: Button, K. David Fulton Publishers, London, pp. 5-47.

some price for this service exceeds the number of seats available – a common occurrence.

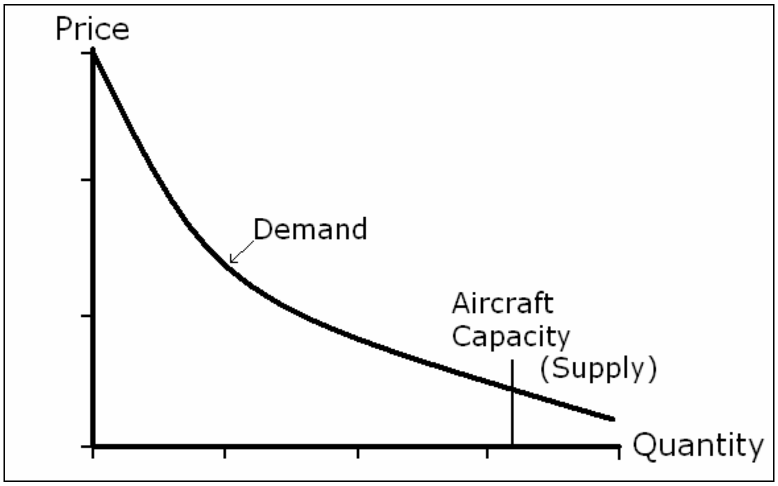

Figure 1: Supply and Demand for Airline Bookings

In the period surrounding deregulation, airlines quickly learned that charging a single fare for flights was not an attractive option in terms of revenue maximization in the newly competitive environment. Consider the hypothetical example where a carrier decides to target high yield business passengers by charging a single high fare throughout the booking process: while yields may be high, planes depart with unused seat inventory because supply exceeds demand at that price. This situation is known as overprotection, because the RM system refuses to sell those empty seats due to its (false) expectation that more high yield demand will materialize. The aircraft ultimately departs with empty seats, which could have been filled by passengers paying lower fares than the carrier was charging. The ratio of passengers to seats, or load factor, is typically low in the case of overprotection.

Conversely, if the carrier seeks to fill its seats all the time, passengers will experience tremendous consumer surplus because the fare most of them pay will be less than their WTP. This situation is known as dilution, as the strong revenue streams an airline expects from its high yield passengers have been diluted with low fares. In this case, the load factor is very high because very few seats will ever go unused. In terms of the supply-demand relationship in Figure 1, these cautionary examples of overprotection and dilution are shown below in Figure 2.

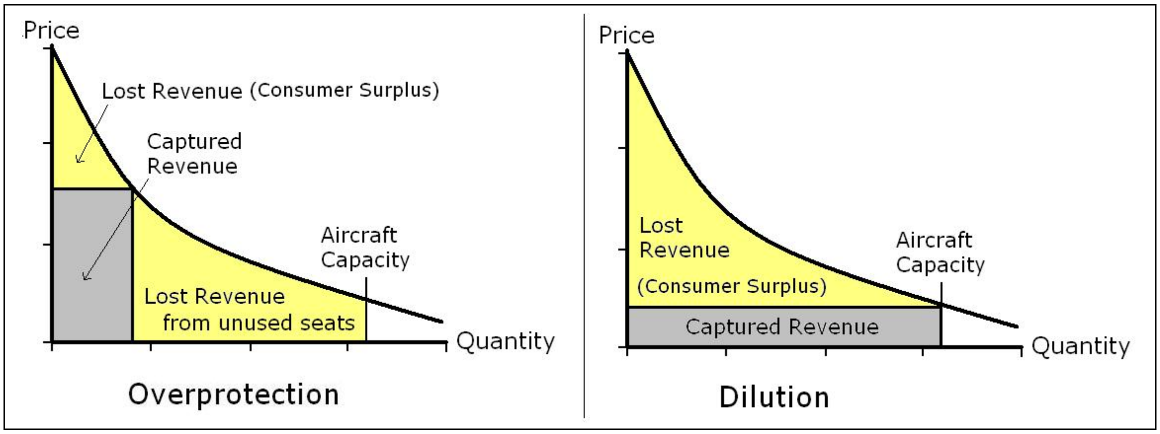

Figure 2: Overprotection and Dilution in Terms of Supply-Demand Relationship

So clearly, single fare policies are unappealing because they result in lost revenue, either from unused seats, diluted fares, or some combination of both. Around the time of deregulation, airline executives and academic researchers alike realized the way to maximize revenue was to charge every passenger his or her maximum WTP – a goal which required different sets of prices for identical units of inventory, or different fares for identical seats on a plane. This concept is known as differential pricing, and represents a core component of successful RM, as described by Belobaba.[6] An illustration of this concept of multiple fares being used to increase revenue in terms of the supply-demand relationship of Figure 1 is shown below in Figure 3.

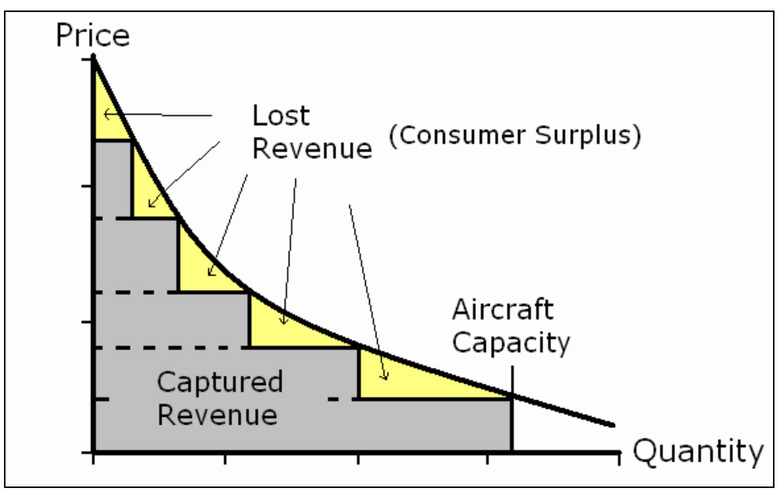

Figure 3: RM's Effective use of Differential Pricing to Counteract Overprotection/Dilution

So when trying to extract maximum revenue from booking passengers, traditional RM makes the classic decision of "whether to take a bird in the hand, or go for two in

[6] Belobaba, P. P. 1998. Airline differential pricing for effective yield management. *The Handbook of Airline Marketing*, D. Jenkins (ed.). The Aviation Weekly Group of the McGraw-Hill Companies, New York, NY, pp. 349-361.

the bush"[7]. The way the field has evolved, today's RM systems optimize the decision of whether or not to sell a ticket at some current price that a potential passenger is willing to pay, or wait for another passenger later in the booking process who will pay more and, thus, contribute more revenue to the airline.

To segregate demand by different WTP levels, traditional differential pricing has attempted to separate business and leisure passengers with "fences" including advance purchase (**AP**) restrictions, mandatory Saturday night stays, round trip purchase requirements, and no-refund or partial-refund policies. The goal of these fences is to add so much disutility to a low fare ticket such that a business passenger, in general, is willing to pay a higher fare (or forced to pay, in the case of advance purchase restrictions) to avoid the restrictions.

1.2 Evolution of the Industry, Growth of "Low Cost Carriers" and Simplified Fare Structures

As the airline industry matured following deregulation, RM systems adapted as well. The hub-and-spoke network model came into widespread use as a way to encourage efficient use of airline resources while connecting countless OD markets which themselves had demand insufficient for direct service. At the same time, RM systems evolved to attempt to account for the network effects of selling not only a common seat inventory to passengers with different WTP levels, but now to passengers with totally different itineraries but who shared a flight leg in the network.

Of course, different airlines embraced (and invested in) RM and operations research in general to different degrees. The decades since deregulation are littered with examples of carriers with sophisticated RM competing directly with carriers less heavily invested in the "new" practice, leading to heavy financial losses, or even bankruptcy for the less adapted airline. As the industry continued to sort itself out during the 1980's and 1990's, a new kind of competitor emerged to challenge the dominant positions of legacy airlines – the low cost carrier (**LCC**).

As the science of RM evolved, it has become common for a single flight leg on a legacy carrier to have dozens of different fares sharing the same seat inventory – with the highest unrestricted business fare being an order of magnitude greater than the lowest, heavily restricted leisure fare. There are countless anecdotal examples of the last minute traveler feeling gouged upon learning the "walk-up" fare he paid was priced at 15 times that for the leisure traveler who booked a month earlier and is sitting in the next seat. And with stories of bad customer service, missed connections, confusing fare structures, and frustrating booking restrictions, legacy carriers undoubtedly left themselves vulnerable in the marketplace.

Enter the LCC. Characterized by more point-to-point service, lower fares, and simplified fare structures, carriers like Southwest, JetBlue, AirTran, and others have eroded the market share of legacy carriers in recent years. And unlike the wave of low cost airlines in the 1980's, which was weathered via RM-enabled fare matching by the legacies, this current incursion of LCCs is having more success thus far against its Major counterparts.

[7] Cook, T. M. 1998. SABRE Soars. *OR/MS Today.* Volume 25, Issue 3, pp. 26-31.

Of course, the troubles of legacy carriers are not entirely due to LCC incursion. The emergence of online airfare searching and booking has removed some of the pricing power legacies traditionally had by instantaneously expanding the consumer marketplace and rewarding diligent searchers with the lowest priced itineraries across the industry. This is to say nothing of overall economic downturns, fuel prices, and air travel demand shocks by the September 11 terrorist attacks, the SARS outbreak in Asia, etc. And, ironically, the confluence of these pressures has led many legacies to reduce costs by way of curtailing comforts like pillows and meals, thus further indistinguishing their on-board product from that of the LCCs – the same LCCs over whom the legacies once justified a price premium with on-board amenities to enhance the flying experience.

Regardless of the causes, LCCs have undoubtedly claimed a significant share of the air transportation market – ECLAT Consulting[8] has pegged it as 18% of 2005 US total air transportation capacity, in terms of Available Seat Miles (ASM) – and legacy carriers have been compelled to respond. Regarding revenue management, several airlines have opted for simpler fare structures in an effort to match the fewer fare classes, relaxed booking restrictions, and fare ceilings trumpeted by LCCs.[9,10,11]

While these new, simplified fare structures may be more appealing to consumers, they present a new set of challenges for RM systems. Because of the fewer fare classes and relaxed booking restrictions, we are faced with a huge problem – no longer can we differentiate a business passenger from her leisure counterpart based on which fare product she chooses.

The classic assumption of demand independence among fare classes – tenuous from the outset – is totally invalid when the consumer is presented with a set of simplified products completely identical save price. And with the internet now enabling market transparency, no longer can we assume that those booking late in the process are business travelers willing to fly at any price.

Clearly, the need exists to rethink the way tickets are sold when no barriers exist separating leisure and business travelers. So just as RM once evolved to fit the changing air transportation marketplace left in deregulation's wake, the field must again adapt to today's changing competitive environment.

1.3 New Approach – Hybrid Forecasting

So the case has been presented for adapting RM in the context of simplified fare structures. These simplifications to traditional fare structures typically consist of relaxations to Saturday night stay, cancellation, refundability, and AP requirements. As discussed previously, these simplifications present a daunting challenge to traditional RM systems – the inability to distinguish between business and leisure demand.

---

[8] ECLAT Consulting. December 2005. Converging Profiles? Presented by Swelbar, W. S. Arlington, VA.
[9] Adams, Marilyn. January 5, 2005. "Delta cuts, simplifies fares." *USA Today*.
[10] Air Canada Press Release. January 26, 2004. "Air Canada introduces low, simplified fares to the United States; Online ticket sales tripled at aircanada.com."
[11] Zellner, W., M. Arndt. December 2, 2002. "American's latest test flight." *Business Week.*

To counteract this challenge, we explore a new methodology – hybrid forecasting – that ignores the conventional business/leisure paradigm in favor of a new distinction – "product-oriented" versus "price-oriented" demand. In short, a product-oriented passenger is only interested in a specific fare product (thus, uninterested in its price), in contrast to the price-oriented passenger whose only objective is to pay the lowest fare possible. Originally developed by Belobaba and Hopperstad[12], the goal of HF is to classify all bookings into one of these two demand categories, and then to predict future demand of both product-oriented and price-oriented passengers in concert.

Because we assume that these two groups of passenger demand exhibit vastly different airline booking behavior, we use different methods of forecasting the demand. Hence, the word "hybrid" can refer to the simultaneous employment of two separate forecasting methods. Another appropriate interpretation of the "hybrid" term in HF could refer to the coexistence of several differentiated fare classes (made unique by booking restrictions) with other undifferentiated fare classes (distinguished only by price).

In conjunction with HF, we also explore the efficacy of two RM tools: Fare Adjustment and Path Categorization. Both of these techniques were designed for network carriers in competitive environments where traditional fare structures have been eroded; thus, their inclusion in this thesis.

1.3.1 Fare Adjustment

This method, developed by Fiig and Isler[13] at Scandinavian Airlines (SAS) and Swissair, respectively, is a technique to augment airline revenues in less restricted fare structure environments. To briefly summarize, FA logic adjusts the fares used by a network seat allocation optimizer – not the actual fares offered to booking passengers – to proactively close selected lower fare classes and induce booking passengers to pay more. Network inventory allocation optimizers are explained further in Section 2.2.3.2, as is FA in Section 2.4.2.

1.3.2 Path Categorization

Intuitively, one would expect that when multiple airlines offer competing service in an OD market, passengers prefer non-stop paths from the Origin to the Destination over connecting service on multiple flight legs, all else being equal. By practicing more aggressive HF (and FA) in its non-stop OD markets, we expect that an airline can enjoy revenue gains via exploiting the higher WTP for direct service. For simplification in this thesis, "Path Categorization" refers to an airline assuming different passenger sell-up behavior in dominant markets, and adjusting its HF and/or FA accordingly to improve revenues.

---

[12] Belobaba, P., C. Hopperstad. 2004. Algorithms for revenue management in unrestricted fare markets. Presented at the Meeting of the INFORMS Section on Revenue Management, Massachusetts Institute oof Technology, Cambridge, MA.
[13] Fiig, T., Isler, K. 2004. "SAS O&D low cost project." *PODS Consortium Meeting*, Minneapolis.

1.4 Objectives of the Thesis

As discussed above in Section 1.2, the simplification of traditional fare structures is common in today's air transportation marketplace, and has the potential to erode revenue by blurring the business/leisure passenger delineation upon which RM evolved. Hybrid forecasting offers a possible solution to this problem by rethinking the crucial question: what kind of demand can be expected for this flight?

Forecasting is the process of quantitatively estimating the expected demand for a particular service, and relies on bookings for previous and current similar services. Needless to say, the demand forecast is a critical component of the RM process; despite its level of precision, a sophisticated (and expensive) RM system can be rendered ineffective when its seat allocation optimizing component is fed a bad forecast. A particular method currently used with some effectiveness is known as pick-up forecasting, and is described in Section 2.2.1.1. However, pick-up forecasting was not developed for simplified fare structures discussed here.

Instead of predicting demand for various (and independent) types of business and leisure passengers, and matching those predictions with a set of customized fare class, HF imagines product-oriented and price-oriented passengers, and separately predicts their numbers throughout the booking cycle, ignoring any notion that a particular passenger is traveling for business or leisure purposes.

Thus, the goal of this thesis is to answer the following question: Does hybrid forecasting lead to revenue improvement over pick-up forecasting when an airline uses a simplified fare structure? Furthermore, this thesis also examines HF in conjunction with Fare Adjustment and/or Path Categorization in an effort to uncover any additional revenue improvements; all performance evaluations employed a competitive airline simulator, as described further in Chapter 3.

1.5 Structure of the Thesis

This thesis is comprised of three parts: a review of the relevant literature, a discussion of the PODS simulator and the approach to HF simulations, and an analysis of the results of those simulations.

Chapter 2 presents a selective discussion of previous work done on revenue management with an emphasis on the problem of simplified fare structures examined in this thesis. Topics covered in the chapter include forecasting, specific RM models, the emergence of simplified fare structures due to LCCs, and a discussion of product-oriented and price-oriented forecasting techniques.

In Chapter 3, we discuss the Passenger Origin-Destination Simulator used for analysis of Hybrid Forecasting. The chapter consists of a general discussion of PODS with a focus on the elements related to HF, as well as an introduction to the specific experiments performed in this thesis.

The detailed methodologies for those simulations as well as their results are presented in Chapters 4 and 5, with an eye on determining revenue improvements enabled by HF. Not only are the potential revenue benefits (or losses) quantified, but we also analyze some of the underlying effects these new RM tools have on loads, yields, fare class mix, etc. in order to isolate the ramifications of each experiment. And Chapter 6 attempts to summarize the experiments performed, as well as the revenue benefits possible with hybrid forecasting alone and with Fare Adjustment and/or Path Categorization; several directions for future work are also presented.

2 Literature Review

This chapter begins by reviewing the use of forecasting in the airline industry, in general, followed by a more focused discussion of Revenue Management (RM) systems, stressing the specific algorithms and methods to be used for simulation. Next, we examine the evolution of the airline industry in general, and the emergence of less restricted fare structures in particular, in an effort to describe the need for changes in conventional RM. Also, two methods recently developed to improve RM in the absence of traditional fare structure restrictions will be presented; Q-forecasting being the price-oriented component of Hybrid Forecasting (HF), and Fare Adjustment (FA) being a supplemental optimization tool to improve network revenues. This chapter concludes with the presentation of price-oriented and product-oriented demand which is the basis of HF.

2.1 Forecasting in the Airline Industry

In his Ph.D. thesis, Lee[14] provides a thorough review of airline passenger forecasting literature to that point; more notably, he also posits that airline demand forecasting can be done on three different levels of varying aggregation: macro, micro, and passenger. And though his own Ph.D. thesis focuses more heavily on demand detruncation, Zeni[15] builds upon Lee's three-tiered classification of airline demand forecasting and provides an overview of forecasting in the industry at large.

Macro-level forecasts represent at the highest aggregation of predicted passenger demand, and can include overall projections of travel demand between two regions during a specific time frame. Taneja[16], de Neufville and Odoni[17], and sections of Sa's[18] Master's thesis provide further discussions and examples of forecasting at the macro-level.

Micro-level forecasts are at a more disaggregate level, such as passenger demand on a specific date or for a specific flight. In this thesis, micro-level forecasting is the most important of the three classifications as it also encompasses the fare class booking estimates essential to effective RM, as described below in Section 2.2. Notable examples and descriptions of micro-level passenger demand forecasts include Ben-Akiva's[19] three models for flight and class specific reservations forecasting and Sa's[18] comparison of time series and regression models for demand by flight and fare class, among many others.

---

[14] Lee, A. O. 1990. Airline reservations forecasting: probabilistic and statistical models of the booking process. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.
[15] Zeni, R. H. 2001. Improved forecast accuracy in revenue management by unconstraining demand estimates from censored data. Ph.D. thesis. Rutgers, the State University of New Jersey, Newark, NJ.
[16] Taneja, N. K. 1978. *Airline traffic forecasting*. Lexington Books, Lexington, MA.
[17] de Neufville, R., A. Odoni. 2003. *Airport systems: planning, design, and management*. McGraw-Hill, New York, NY.
[18] Sa, J. 1987. Reservations forecasting in airline yield management. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.
[19] Ben-Akiva, M. 1987. Improving airline passenger forecasts using reservation data. Paper presented at the Fall ORSA/TIMS Conference, St. Louis, MO.

And finally, passenger level forecasting represents the highest resolution of demand modeling – the choices facing individual passengers; examples include models for airline choice, airport choice, itinerary choice, and even fare class choice. Further examples and/or models at the individual traveler level are given by Kanafani[20], Ben-Akiva and Lerman[21], and Boeing[22] (as discussed below in Section 3.1).

2.2 Revenue Management Component Models

There is no shortage of literature on revenue management. McGill and van Ryzin[2] provide a summary of RM's history along with descriptions of many innovations over the years, as does Talluri and van Ryzin's[23] book, though in a more technical fashion. Boyd and Bilegan[24] present a well organized and more technically detailed survey of many of the same developments, but not limited to the context of airline RM. Conversely, Clarke and Smith[25] limit their scope to the airline industry, but discuss a range of operations research contributions, of which RM is included.

Similar to Clarke and Smith, Barnhart et al.[26] present more detailed overviews of fleet assignment, revenue management, and aviation infrastructure operations. Regarding the evolution of RM, they discuss the first generation systems of the early 1980's which simply collected and stored reservation data from a computer reservations system; this process could be more aptly described as data collection or revenue observation than actual revenue management. The second generation of RM systems in the mid-1980s allowed airlines to follow bookings prior to a flight's departure and compare to expected booking patterns.

Operations research advances were finally integrated into the RM process with the third generation RM systems developed in the late 1980's and early 1990's. As shown below in Figure 4, a typical third generation system took separate pricing structures (revenue data), the airline's database of historical bookings, current booking data, as well as cancellation and no-show data and fed it into three component models: the demand forecaster, the overbooking module, and the fare class mix (seat allocation) optimizer. The output generated from this process was the optimal booking limits for each flight and fare class. According to Barnhart et al., most medium and large-sized carriers worldwide have implemented third generation RM systems with architecture similar to that of Figure 4, and typically they enjoy revenue gains of 4%-6% compared to no seat inventory control.

This section of the literature review focuses on the three modeling components of a third generation RM system: forecasting, overbooking, and seat allocation. Special emphasis is given to forecasting as it is the subject of this thesis.

---

[20] Kanafani, A. K. 1983. *Transportation demand analysis*. McGraw-Hill, New York, NY.
[21] Ben-Akiva M., S. Lerman S. 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand*. MIT Press, Cambridge, MA.
[22] Boeing Airplane Company. 1997. *Decision Window Path Preference Methodology Description.* Seattle, WA.
[23] Talluri, K., G. van Ryzin. 2004. *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers, Dordrect, Netherlands.
[24] Boyd, E. A., Bilegan, I. C. 2003. Revenue management and e-commerce. *Management Science*. Volume 49, Issue 10, pp. 1363–1386.
[25] Clarke, M., B. Smith. 2004. Impact of operations research on the evolution of the airline industry. *Journal of Aircraft*. Volume 41, Issue 1, pp. 62-72.
[26] Barnhart, C., P. Belobaba, A. R. Odoni. 2003. Applications of operations research in the air transport industry. *Transportation Science*, Volume 37, Issue 4, pp. 368-391.

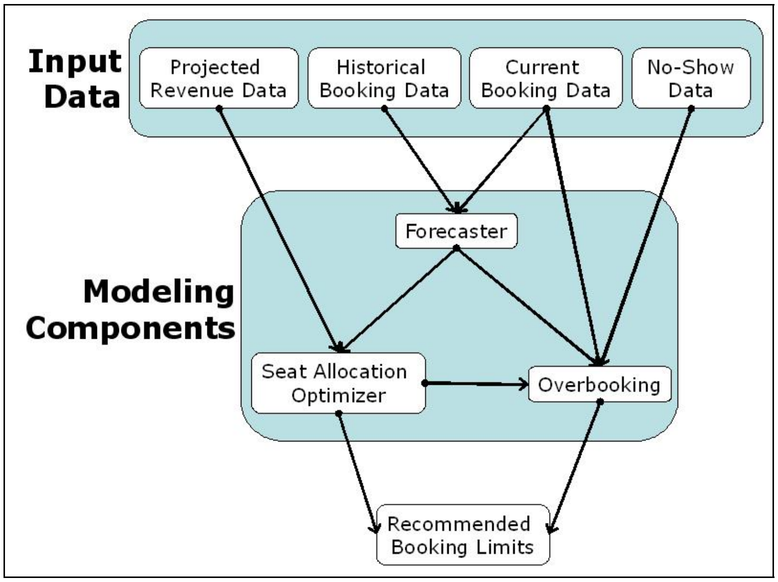

Figure 4: "Third Generation" RM System (Barnhart et al.[26])

2.2.1 Traditional Forecasting Models

As mentioned in Section 2.1, the forecaster in an RM system works at a micro-level, generating demand forecasts by fare class, either for individual flight legs or for overall Origin to Destination (OD) itineraries. Weatherford[27] and Zeni[15] refer to five models commonly used in practice, which are mentioned in the sections below. Because HF makes use of pick-up forecasting as one of its components, that particular model is reviewed in greater depth in the next section.

2.2.1.1 Pick-up Forecasting

A pick-up model of demand is a simple forecasting technique that has proven to be effective under the traditional assumptions of RM. In this method, the "pick-up" can be described as the forecasted number of incremental bookings over a specified future time period based upon historical trends. The pick-up is generally added to the number of current bookings to forecast the total demand at the end of the specified period.

There are actually two versions of the pick-up model and they are similar in formulation: the classical pick-up model and the advanced pick-up model. The classical model uses only data from departed flights, while the advanced model (developed by L'Heureux[28] at Canadian Pacific) also makes use of data from flights that have not yet departed; only classical pick-up was used in the simulations for this thesis.

---

[27] Weatherford, L. 1999. Forecast aggregation and disaggregation. *IATA Revenue Management Conference Proceedings*.

[28] L'Heureux, E. 1986. A new twist in forecasting short-term passenger pickup. *26th AGIFORS Annual Symposium Proceedings*, Bowness-on-Windemere, England, pp. 248–261.

For more information on the pick-up forecasting model, the reader is referred to Zickus[29], Skwarek[30], Usman[31], Gorin[32], or Wickham[33].

2.2.1.2 Other Forecasting Methods

Besides the pick-up model, some other common RM forecasting methods include exponential smoothing, moving average, regression, and multiplicative pick-up[27]. Exponential smoothing (a form of time-series forecasting model), the moving average method, and the multiplicative pick-up model are discussed by Zeni[15]. And Zickus[29], Skwarek[30], Usman[31], Gorin[32], and Wickham[33] all discuss versions of regression forecasting.

In particular, Wickham's Master's thesis includes a comparison of the relative performance of several forecasting techniques, including simple time-series, linear regression, and pick-up models. In general, his tests found that pick-up models consistently outperformed the other methods by leading to the largest revenue contributions.

2.2.2 Overbooking Models

Perhaps the earliest component of RM to be developed, overbooking models were originally used to intentionally accept more reservations than seats existed for flights in an effort to reduce the revenue loss and seat spoilage from no-shows and cancellations. Beckman[34] produced and early, static overbooking model, while other advances were made by Taylor[35], Simon[36], Rothstein[37,3], and Vickrey[38]. For a more thorough literature review of overbooking, the reader is referred to McGill and van Ryzin[2]. No overbooking models were used in the simulations for this thesis.

2.2.3 Seat Allocation Optimizers

---

29 Zickus, J. S. 1998. Forecasting for airline network revenue management: revenue and competitive impacts. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.
30 Skwarek, D. K. 1996. Competitive impacts of yield management systems components: forecasting and sell-up models. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.
31 Usman, A. S. 2003. Demand forecasting accuracy in airline revenue management: analysis of practical issues with forecast error reduction. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.
32 Gorin, T. O. 2000. Airline revenue management: sell-up and forecasting algorithms. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.
33 Wickham, R. R. 1995. Evaluation of forecasting techniques for short-term demand of air transportation. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.
34 Beckman, J. M. 1958. Decision and team problems in airline reservations. *Econometrica*. Volume 26, pp. 134-145.
35 Taylor, C. J. 1962. The determination of passenger booking levels. *2nd AGIFORS Annual Symposium Proceedings*, Fregene, Italy.
36 Simon, J. L. 1968. An almost practical solution to airline overbooking. *Journal of Transport Economics and Policy*. Volume 2, pp. 201–202.
37 Rothstein, M. 1968. Stochastic models for airline booking policies. Ph.D. thesis, Graduate School of Engineering and Science, New York University, New York, NY.
38 Vickrey, W. S. 1972. Airline overbooking: some further solutions. *Journal of Transport Economics and Policy*. Volume 6, pp. 257-270.

Typically thought of as the essence of RM, it is the seat allocation optimizer component of the RM system that actually sets the booking limits in an effort to maximize the revenue contribution of every passenger booking with the airline. It is here that we make the important distinction between Fare Class Yield Management algorithms and OD algorithms based upon the level of optimization (flight leg versus OD path).

2.2.3.1 Fare Class Yield Management – Leg-based Control

Utilizing demand forecasts for individual flight legs, fare class yield management (FCYM) systems use optimizers which determine seat allocation for the set of fare classes on each leg within a network. The most commonly used fare class mix allocation is the idea of serial "nesting" of fare classes – a problem first solved by Littlewood[39] at BOAC for the case of two fare classes. As opposed to allocating seats in partitioned fare classes, nesting instead protects seats in high fare classes by limiting the number of seats sold in lower fare classes based on a forecast of demand for each class, as well as the expected seat revenue.

Belobaba extended the nested seat allocation problem from Littlewood's two classes to multiple fare classes with the Expected Marginal Seat Revenue (**EMSR**) heuristic in his Ph.D. thesis[40]. This algorithm employs leg-based demand forecasts by fare class to produce leg-based seat protection levels for nested booking classes.

EMSR determines booking limits based upon the expected marginal revenue – the probability of selling an additional seat in a given fare class multiplied by the revenue gained from selling that seat. As the number of seats protected in a particular fare class increases, the probability of selling that next seat decreases; thus, the booking limit for a fare class is determined when the EMSR is equal to the fare of the next lower class.

Belobaba's updated EMSRb[41] algorithm protects joint upper classes from the next fare class just below, and has become somewhat of an industry standard for establishing booking limits on a flight leg basis. In order to simplify calculations for joint classes, the fare class demands are assumed normal and independent – assumptions which are violated when the fare class restrictions are eased, as discussed below in Section 2.3. More information on the EMSRb algorithm can be found in Lee[42] and Williamson[43].

In an alternative approach to the multiple nested class problem, optimal formulations for leg/class seat allocation have been (independently) introduced by Brumelle and

---

[39] Littlewood, K. 1972. Forecasting and control of passenger bookings. *12th AGIFORS Annual Symposium Proceedings*, Nathanya, Israel, pp. 95–117.
[40] Belobaba, P. P. 1987. Air travel demand and airline seat inventory management. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.
[41] Belobaba, P. P. 1992. "Optimal versus heuristic methods for nested seat allocation." *AGIFORS Reservations Control Study Group Meeting*. Brussels, Belgium.
[42] Lee, A. Y. 1998. Investigation of competitive impacts of origin-destination control using PODS. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.
[43] Williamson, E. L. 1992. Airline network seat inventory control: methodologies and revenue impacts. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.

McGill[44], Curry[45], Robinson[46], and Wollmer[47], though they require much more significant computational effort and have been shown to only marginally outperform the EMSRb heuristic in terms of revenue gained[24].

2.2.3.2 Network OD Control Models

The major shortcoming of fare class RM algorithms within a network carrier is their failure to model booking behavior in terms of the actual OD products passengers purchase, and instead simplifying analysis to a flight leg basis. In general, we define a path as a feasible set of flight legs (or just one leg for non-stop services) connecting an Origin with a Destination within a given network.

By ignoring network effects, leg-based inventory optimizers run a high likelihood of giving preference to local passengers on a given flight leg at the expense of higher fare connecting passengers on OD paths which include that particular leg. For this reason, much effort has been expended to develop algorithms for path-based protection of booking classes (or OD control)[48].

The use of "virtual buckets" to compare network value of local and connecting fare classes is one approach to network OD control, as described by Vinod[49]. The specific method known as Displacement Adjusted Virtual Nesting (**DAVN**), described by Williamson[43], was used for simulation in this thesis. DAVN couples OD forecasting with leg-based seat inventory control, and uses a deterministic linear program (LP) with an objective of network revenue maximization to calculate a "pseudo fare" for each fare class in the network; this pseudo fare corrects the regular fare for network displacement effects. By grouping each leg's pseudo fares into similar sets, or buckets, and then optimizing booking limits (in a manner similar to EMSRb) for those buckets, the airline has a mechanism for maximizing revenue while accounting for displacement costs over its network. Lee[42] and Williamson[50] both discuss virtual bucketing, including DAVN.

Another approach developed for OD control was the use of bid prices[48,51]. Specific bid price algorithms include the Network Bid Price (NetBP) method, the Heuristic Bid

---

[44] Brumelle, S. L. and McGill, J. I. 1988. Airline seat allocation with multiple nested fare classes. Paper presented at the Fall ORSA/TIMS Conference, Denver, CO. Also presented at the University of British Columbia, 1987.
[45] Curry, R. E. 1990. Optimal airline seat allocation with fare classes nested by origin and destinations. *Transportation Science*. Volume 24, Issue 3, pp. 193–204.
[46] Robinson, L. W. 1995. Optimal and approximate control policies for airline booking with sequential nonmonotonic fare classes. *Operations Research*. Volume 43, Issue 2, pp. 252–263.
[47] Wollmer, R. D. 1992. An airline seat management model for a single leg route when lower fare classes book first. *Operations Research.* Volume 40, Issue 1, pp. 26–37.
[48] Smith, B. C., C. W. Penn. 1988. Analysis of alternative origin-destination control strategies. *28th AGIFORS Annual Symposium Proceedings*, New Seabury, MA.
[49] Vinod, B. 1995. Origin and Destination Yield Management. *The Handbook of Airline Economics*, D. Jenkins (ed.). The Aviation Weekly Group of the McGraw-Hill Companies, New York, NY, pp. 459-468.
[50] Williamson, E. L. 1988. Comparison of optimization techniques for origin-destination seat inventory control. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.
[51] Simpson, R. W. 1989. Using network flow techniques to find shadow prices for market and seat inventory control. Memorandum M89-1, MIT Flight Transportation Laboratory, Cambridge, MA.

Price (HBP) method developed by Belobaba[52], and the Prorated Bid Price (ProBP) as described by Bratu[53].

2.3 The Emergence of Low Cost Carriers and Simplified Fare Structures

As described in the previous section, the field of airline revenue management has grown quite complex in a short amount of time. Since the CAB's deregulation of the airline industry, there have been tremendous advancements in the science of RM, including continuing improvements in forecasting, overbooking, and inventory optimization models (including OD control). This section describes a further evolution of the airline industry – the low cost carrier (LCC) – as well as the growing use of less restricted fare structures and the need to adapt traditional RM systems to handle these changes.

2.3.1 Characteristics of LCCs

There is no one watershed event which marked the inception of low cost carriers. Even though it is widely accepted that 1978's deregulation was the impetus for widespread change throughout the industry[5], LCC archetype Southwest Airlines began operations years earlier in 1971[54], and numerous new-entrant airlines have come and gone in the years since deregulation. In his Ph.D. thesis, Gorin[55] provides a comprehensive summary of changes in the U.S. airline industry since deregulation, focusing on low-fare new entrant airlines in an attempt to characterize and define the LCC.

But just as no distinct impetus exists for the low cost carriers' development, no single definition of LCC (or legacy carrier, for that matter) suffices in today's industry. Supplementing Gorin's[55] description of the LCC business model, Dunleavy and Westerman[56] contrast LCCs with legacies, as do Weber and Thiel[57], among others; Table 1 below provides a brief comparison of some selected characteristics traditionally associated with legacies and LCCs.

---

[52] Belobaba, P. P. 1998. The evolution of airline yield management: fare class to origin-destination seat inventory control. *The Handbook of Airline Marketing*, D. Jenkins (ed.). The Aviation Weekly Group of the McGraw-Hill Companies, New York, NY, pp. 285-302.
[53] Bratu, S. J-C. 1998. Network value concept in airline revenue management. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.
[54] Gittell, J. H. 2003. *The Southwest Airlines way: using the power of relationships to achieve high performance.* McGraw-Hill, New York, NY.
[55] Gorin, T. O. 2004. Assessing low-fare entry in airline markets: impacts of revenue management and network flows. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.
[56] Dunleavy, H., Westermann, D. 2005. Future of airline revenue management. *Journal of Revenue and Pricing Management*. Volume 3, Issue 4, pp. 380-282.
[57] Weber, K., Thiel, R. 2004. "Optimisation issues in low cost revenue management." *AGIFORS Reservations & Revenue Management Study Group Meeting*, Auckland, New Zealand.

Table 1: Differences Between Traditional Legacy Carrier and LCC Business Models (Gorin[55], Dunleavy and Westerman[56], Weber and Theil[57])

Legacy Carriers	Low Cost Carriers
• Mix of low and high fare products	• Typically low fares
• Up to dozens of fare products per OD pair	• ~5 fare products per OD pair
• Distribution via travel agents, third party vendors, carrier's website…	• Direct distribution and simplified passenger processing
• Low aircraft utilization	• High aircraft utilization
• Low labor productivity	• High labor productivity
• Round trip & one-way fares	• Only one-way fares
• Typically connecting traffic through hubs	• Typically point-to-point traffic
• Complex segmentation through fare structures (business, leisure, etc.)	• No customer segmentation
• On-board amenities (meals, pillows, etc.)	• "No frills"
• Variety of aircraft types in fleet	• Very few aircraft types

However, it is important to note that the legacy and LCC characteristics in Table 1 are generalities and not exclusive to either column. For example, many of the larger legacy carriers have made significant efforts to reduce costs by reducing the number of aircraft types in the fleet and eliminating on-board amenities. Conversely, several LCCs operate large hub operations and offer amenities like in-seat television.

Ironically, the dramatic cost reductions undertaken at the legacies, coupled with the gradually increasing cost structures at several LCCs have led some to question if "low cost carriers" will continue to be characterized by lower costs. In fact, ECLAT Consulting[8,58] has proposed branding this segment of the industry with the new moniker "large market oriented carriers" (LMOs) due to their propensity for point-to-point service between major population centers. However, an exhaustive discussion on the growing similarities between LCCs and legacy carriers is beyond the scope of this thesis, and this particular market segment will continue to be referred to as "LCCs" for the remainder of this work.

In his Ph.D. thesis, Gorin[55] examines the impact LCCs have had on the industry's traditional giants, concluding that gradually increasing competition has contributed to the weakening of the legacy carriers. In a presentation for the MIT Global Airline Industry Program, Swelbar[59] details the extent to which LCCs have penetrated the U.S air transportation network and demonstrates that legacy carriers may have been experiencing profit deterioration since 1998. And in an informational brief filed along with its December 2002 bankruptcy, United Airlines[60] seemingly concedes that revenues eroded due to added competition from LCCs, critically crippling the legacy carrier to the point that drastic changes were sorely needed.

[58] ECLAT Consulting. 2005. Repealing the Wright Amendment – risks facing small communities and the Dallas Metroplex. Arlington, VA.
[59] Swelbar, W. S. 2002. The role and impact of low cost carriers. *1st Annual MIT Airline Industry Conference*, Washington, DC. ECLAT Consulting, Arlington, VA.
[60] United Airlines. December 9, 2002. Informational Brief of United Airlines, Inc. Filed in the Bankruptcy Court for the Northern District of Illinois Eastern Division.

Tretheway[61] believes that the traditional legacy carrier business model is dangerously flawed, and Franke[62] analyzes the adverse financial effects LCCs have had on the legacies, while arguing for an overhaul of the traditional network carrier business model.

In this new competitive environment, Ratliff and Vinod[63] explain how legacy carriers often are compelled to match the fares – and sometimes even entire fare structures – of their LCC counterparts. As mentioned previously, traditional network carriers such as Delta[9], Air Canada[10], and American[11], among others, have experimented with the simpler fare structures of LCCs in an attempt to remain competitive. And Donnelly et al. [64] describe the RM challenges faced by bmi when the carrier simplified its fare structures to one-way, restriction-free products due to competitive pressures from European LCCs in 2002.

2.3.2 The Inadequacy of RM Systems under Fare Simplifications

In a largely unscientific and anecdotal analysis, Kuhlmann[65] asserts that the traditional revenue management systems have been rendered ineffective for three reasons: the growth of low cost carriers (as discussed above), increased pricing transparency for consumers due to the internet, and a general downturn in the economy beginning somewhere near the year 2000. He continues by arguing that management teams at legacy airlines in the late 1990's were largely oblivious to the emerging LCCs, unresponsive to customer dissatisfaction, and unconcerned with investing in RM systems suited for LCC competition. In rebuttal, Cary[66] refutes Kuhlmann's criticisms of legacy management, yet concedes the fact that traditional RM systems are ill suited for environments where "business customers are demonstrating unusual levels of price-sensitivity." So in essence, RM systems are unable to distinguish between business and leisure travelers. And when the conventional fare structures are dismantled, the restrictions fencing business and leisure travelers into their respective classes disappear, making it impossible for an RM system to effectively segment demand.

And according to Boyd and Billegan[24], the common assumption of demand independence is questionable even with the traditional fare structures and their bevy of booking restrictions. Absent those restrictions, argue Ratliff and Vinod[63], this assumption is totally invalid as no fences exist to prevent business passengers from booking in a fare class priced below his/her willingness-to-pay (WTP).

Thus, when fare structures are simplified, conventional RM systems used by legacy carriers become inadequate for the two primary reasons described above:

---

[61] Tretheway, M. W. 2004. Distortions of airline revenues: why the network airline business model is broken. *Journal of Air Transport Management*. Volume 10, Issue 1, pp. 3-14.
[62] Franke, M. 2004. Competition between network carriers and low-cost carriers – retreat battle or breakthrough to a new level of efficiency?. *Journal of Air Transport Management*. Volume 10, Issue 1, pp. 15-24.
[63] Ratliff, R., Vinod, B. 2005. Airline pricing and revenue management: a future outlook. *Journal of Revenue and Pricing Management*. Volume 4, Issue 3, pp. 302-307.
[64] Donnelly, S., James, A., Binnion, C. 2004. bmi's response to the changing European airline marketplace. *Journal of Revenue and Pricing Management*. Volume 3, Issue 1, pp. 10-17.
[65] Kuhlmann, R. 2004. Why is revenue management not working?. *Journal of Revenue and Pricing Management*. Volume 2, Issue 4, pp. 378-387.
[66] Cary, D. 2004. A view from the inside. *Journal of Revenue and Pricing Management*. Volume 3, Issue 2, pp. 200-203.

1. The erosion of customer segmentation practices, both due to the elimination of fare class restrictions and a growing unwillingness of business travelers to pay fares so much higher than their leisure counterparts.
2. With no differentiator among fare classes except price, passengers naturally buy the lowest class available. Besides clearly violating the demand independence among fare classes assumption, it also becomes difficult for a RM system to generate accurate demand forecasts due to the dearth of bookings in higher fare classes.

The direct consequence of this fare structure simplification described above – passengers booking below their WTPs – is known as "buy-down", and undoubtedly leads to revenue dilution, as described above in Section 1.1.

A more indirect consequence of easing fare structure restrictions is known as "spiral-down", and is depicted below in Figure 5. Due to the structure of RM systems, future booking forecasts are directly dependent upon historical bookings, as described above in Section 2.2.1. Due to buy-down, the RM system records fewer bookings in higher fare classes. As a result, the forecaster produces a lower projection of high-fare demand, which leads the optimizer to protect fewer seats in the higher fare classes and make more seats available for the lower fare classes. Of course, the surplus of lower fare class seats starts the cycle over again, with revenues becoming more diluted with each iteration.

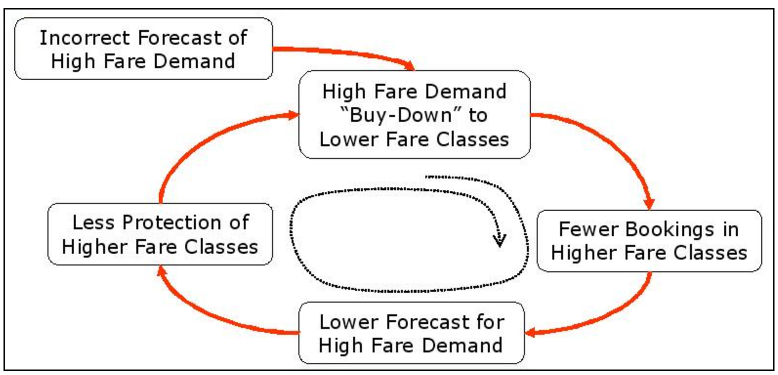

Figure 5: Spiral-Down Effect

A result of making incorrect assumptions about customer booking behavior, the spiral-down effect has been discussed by both Kleywegt et al.[67] and Cooper et al.[68] with an aim of modeling more general violations of traditional RM's assumptions (i.e. not for the specific case of simplified fare structures discussed above).

Both the buy-down and spiral-down effects as a result of fare structure simplifications have been described and analyzed in practical application at United

[67] Kleywegt, A. J., T. Homem-de-Mello, W. L. Cooper. 2003. Models of the spiral down phenomenon. Paper presented at the Meeting of the INFORMS Section on Revenue Management, Columbia University, New York, NY.
[68] Cooper, W. L., T. Homem-de-Mello, A. J. Kleywegt. 2004. Models of the spiral-down effect in revenue management. Working Paper, Department of Mechanical Engineering, University of Minnesota, Minneapolis, MN.

Airlines by Ozdaryal and Saranathan[69] and in the Passenger Origin-Destination Simulator by Cusano[70] and Cléaz-Savoyen[71].

2.4 RM tools for the New Competitive Environment

2.4.1 Q-forecasting

As described above, legacy carriers often react (sometimes out of necessity) to competition with LCCs by significantly changing their fare structures and loosening booking restrictions and/or advance purchase requirements. But in doing so, traditional RM systems become ineffective due to the violation of the demand independence assumption among fare classes – a violation which invalidates the use of established demand forecasting techniques.

To deal with totally unrestricted fare structures (that is, the only differentiator among fare classes is price), Belobaba and Hopperstad[12] developed "Q-forecasting." This forecasting method seeks to forecast demand only in the lowest class (denoted as Q-class) and then uses estimates of passenger WTP to close lower fare classes and force "sell-up" into higher ones, as shown below in Figure 6.

By transforming all historical bookings (regardless of fare class) into an equivalent number of Q-bookings, the Q-forecaster can estimate the number of bookings possible in each fare class assuming a certain level of passenger WTP. Thus, the inventory optimizer can maximize revenue by strategically closing lower classes and forcing a certain fraction of the Q-bookings to sell up into higher fare classes. In his Master's thesis, Cléaz-Savoyen[71] examined the process and concluded that Q-forecasting was an effective technique for forecasting when restriction-free fare structures are used.

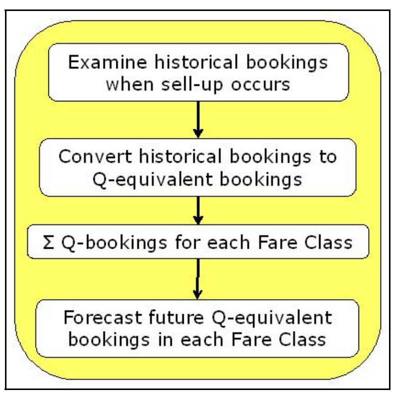

Figure 6: Basic Q-forecasting Logic (Cléaz-Savoyen[71])

---

[69] Ozdaryal, K., Saranathan, B. 2004. "Revenue management in broken fare fence environment." *AGIFORS Reservations & Revenue Management Study Group Meeting*, Auckland, New Zealand.
[70] Cusano, A. J. 2003. Airline revenue management under alternative fare structures. Master's thesis, Massachusetts Institue of Technology, Cambridge, MA.
[71] Cléaz-Savoyen, R. L. 2005. Airline revenue management for less restricted fare structures. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

2.4.2 Fare Adjustment

As described and tested by Cléaz-Savoyen[71], Fare Adjustment (FA) is a method of reconciling the coexistence of different fare structures sharing the same flight leg. In his thesis, Cléaz-Savoyen tested FA in DAVN as a method of having a traditional network carrier's restricted fare structure share virtual buckets with an unrestricted fare structure brought about by the entrance of an LCC to certain OD markets. The conflict arises because of the different characteristics of the two structures. Demand in the traditional restricted structure is assumed to be segmented, as discussed above. However, in the unrestricted fare structure, there is no segmentation of demand and revenues are only maximized by selling-up passengers to higher fare classes after closing lower classes.

So for an airline using DAVN, it may come to be that the two OD fare classes coexist in the same bucket – one for the restricted fare structure and the other for the unrestricted one. It is possible that the closure of that particular fare bucket under DAVN logic may be the optimal strategy for the one structure, but suboptimal for the other. He illustrates one such conflict below in Figure 7.

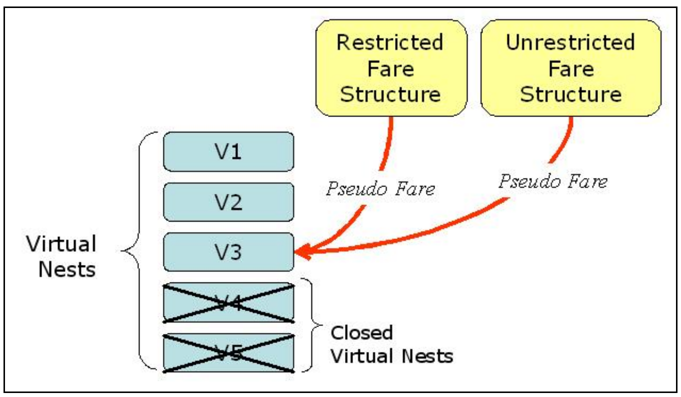

Figure 7: Fare Classes and Virtual Buckets Before FA (Cléaz-Savoyen[71])

Using the Fare Adjustment techniques developed by Fiig and Isler[13], the pseudo fare for the unrestricted fare structure (not the actual fare offered to passengers) can be lowered by a certain amount in order to shift it to a lower virtual bucket, and thus, close that unrestricted fare class sooner. The quantity of this extra decrease is referred to as the Price Elasticity cost ("PE cost"), as shown below in Figure 8, and depends on estimates of passenger WTP, as with Q-forecasting. This PE cost accounts for the risk of buy-down in the unrestricted fare structure. In theory, both fare structures are will act more independently, and revenue gains will be realized network wide.

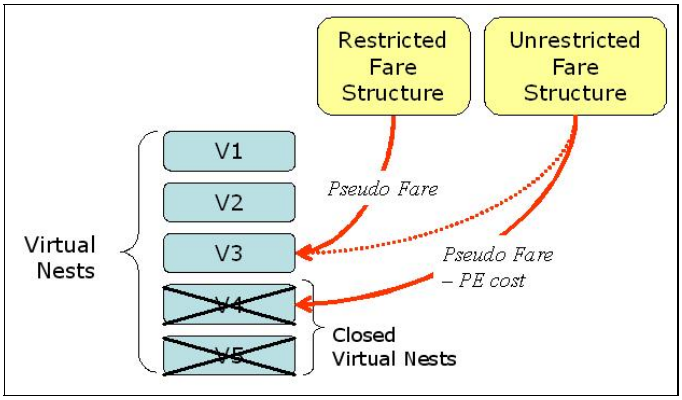

Figure 8: Fare Classes and Virtual Buckets After FA (Cléaz-Savoyen[71])

Cléaz-Savoyen[71] tested the FA methodology in the context of the DAVN optimization process and concluded that FA had the potential to be an effective technique for forecasting in restriction-free fare structure environments. In this thesis, we test its applicability in simplified fare structure environments, which share the characteristic of multiple undifferentiated fare classes of the unrestricted fare structures.

2.5 Price-oriented versus Product-oriented Demand

For all the attention given to those restriction-free fare structures in the development of techniques like Q-forecasting and Fare Adjustment, it is rarely the case in today's industry where any airline – LCC or legacy carrier – uses a fare structure with multiple products differentiated by price and price alone. Airlines are much more likely to use a so-called simplified fare structure (also referred to as semi-restricted) than a totally unrestricted one. In a semi-restricted fare structure, there often exist several lower fare classes that are undifferentiated (expect by price, of course) just below two or three higher fare classes which are differentiated by a restriction (or lack thereof) uncommon to those below.

In this case, a method like Q-forecasting is not completely appropriate because the fare structure is not totally unrestricted, and the differentiation between fare products will keep certain passengers from buying in certain classes for reasons other than price. But traditional RM forecasting techniques will also prove suboptimal because of the presence of those undifferentiated fare classes which, again, invalidate the assumption of independence among classes. Clearly, the demand must be modeled as some mix of passengers who are price sensitive and those who are shopping for a specific product.

According to Boyd and Kallesen[72], the changing business environment has made irrelevant the traditional demand segmentation which characterized RM systems. The increased ease and transparency of online booking and the less restricted fare structures being introduced to the industry at large by LCCs has led to more advanced traveler trip purpose anonymity. No longer is it possible to recognize a passenger as either business or leisure based solely upon booking behavior. In a departure from that typical business versus leisure passenger mix, Boyd and

[72] Boyd, E. A., Kallesen, R. 2004. The science of revenue management when passengers purchase the lowest available fare. *Journal of Revenue and Pricing Management*. Volume 3, Issue 2, pp. 171-177.

Kallesen propose a new segregation of demand that appears to suit semi-restricted fare structures quite well: yieldable versus priceable passengers.

2.5.1 Product-oriented Demand (Yieldable)

Under a yieldable model of demand, a product-oriented passenger is just that – specifically interested in a particular product, and only that particular product. Under this model of demand, our previous assumption of demand independence among fare classes is again valid because each fare class product is matched with a specific product-oriented set of passengers who are uninterested in other fare classes.

In terms of the implications for forecasting, the traditional method of pick-up forecasting remains relevant for predicting future demand of product-oriented passengers (and product-oriented passengers only) as it has always relied on independence among fare classes, as discussed in Section 2.2.1.

In short, using a product-oriented model of demand changes nothing but terminology for our traditional RM systems – it has been used all along. Thus, the literature is quite extensive on the subject of yieldable demand, as it encompasses nearly every major development made in revenue management; as such, Boyd and Kallesen suggest Boyd and Bilegan[24] for a review of product-oriented demand.

2.5.2 Price-oriented Demand (Priceable)

Totally opposite from product-oriented demand, an "idealized" price-oriented passenger has no concept of the characteristics of multiple fare products and is simply interested in purchasing airfare at the lowest price, regardless of his or her WTP. The notion of priceable demand has come to the forefront with all the work on restriction-free fare structures. After all, if there are absolutely no differentiating characteristics among fare classes, every rational passenger exhibits priceable behavior because he or she will purchase at the lowest price, possibly unaware that multiple products are actually being sold.

Because of the generalized nature of the optimal booking limit methods proposed by Brumelle and McGill[44], Curry[45], Robinson[46], and Wollmer[47], Boyd and Kalleson suggest that priceable demand can be analyzed "in a manner similar in spirit to that used by" those researchers, however noting that the price-oriented demand "leads to a much more complicated analysis."

Regarding the use of these optimal booking limit methods for priceable demand, even Curry acknowledges the computational difficulty in such an approach, writing "Optimal allocations are much more complex if class demands are not independent…An iterative approach may be possible, but seems unwieldy and unnecessary."

Regarding forecasting for price-oriented demand, Belobaba and Hopperstad's[12] Q-forecasting technique described in Section 2.4.1 is valid because it does not make the classic assumption of demand independence. Furthermore, it was designed for the environment of totally unrestricted fares, so it is seemingly well suited for passengers to which price is the only factor.

2.5.3 Product-oriented and Price-oriented Demand Together

As described above, passenger demand for air transportation can be modeled as a mix of product-oriented and price-oriented demand, and a separate demand forecasting method is needed for each segment. But when the wrong forecaster is used, the seat inventory optimizer will produce suboptimal results.

According to Boyd and Kallesen, if a forecasting model for product-oriented demand (like pick-up forecasting, as discussed in Section 2.5.1) is used to project price-oriented demand, the forecast will undoubtedly overestimate demand for the lower fare classes and authorize too many low-yield seats leading to both the (immediate) buy-down effect and the (eventual) spiral-down effect discussed in Section 2.3.2. So the example of revenue loss due to dilution presented in Figure 2 will be realized.

On the other hand, if a forecasting model for price-oriented demand (like Q-forecasting, as discussed in Section 2.5.2) is used to project product-oriented demand, the forecast will overestimate demand for the higher fare classes and protect too many high-yield seats leading to revenue loss as those seats go unnecessarily unsold. So the example of revenue loss due to overprotection also presented in Figure 2 above will be realized.

2.5.4 Summary of Product-oriented and Price-oriented Demand

We can summarize Section 2.5 as follows:

- Demand is actually a mix of product-oriented and price-oriented passengers;
- Pick-up forecasting can be used to predict product-oriented demand;
- Q-forecasting can be used to predict price-oriented demand;
- The use of pick-up forecasting to predict price-oriented demand should be minimized in order to avoid revenue dilution;
- The use of Q-forecasting to predict product-oriented demand should be minimized in order to avoid revenue loss due to overprotection;

2.6 Chapter Summary: The Need for “Hybrid” Forecasting

We began this chapter with a review of passenger demand forecasting and its use within the airline industry at large; note that this thesis focuses on micro-level demand forecasts used to predict future bookings for Revenue Management systems. In Section 2.2 we reviewed the extensive literature on RM, discussing a typical RM system and its three component models: forecasting, overbooking, and optimizing seat allocation.

We then turned our attention to Low Cost Carriers and their role in expanding the use of simplified fare structures in Section 2.3; this section continued by describing the inadequacy of traditional RM systems in these less restricted fare structures. And in Section 2.4 we described two new RM developments designed for use in these simplified structures: Q-forecasting and Fare Adjustment.

Finally in Section 2.5, we discussed the idea of modeling overall passenger demand as the sum of two contrasting components: product-oriented and price-oriented demand.

Techniques exist for forecasting both product-oriented and price-oriented demand separately. But to effectively practice RM in a semi-restricted fare structure environment, we should use a "hybrid" forecaster which combines product-oriented and price-oriented demand forecasting. This thesis will present several methods of classifying product-oriented and price-oriented demand in order to define a successful hybrid forecasting approach; this thesis will then present several situations in which HF can be improved in terms of revenue gained in order to address the challenges posed by RM systems in today's airline industry.

3 The PODS Approach to Revenue Management Simulation

In the context of airline revenue management (RM), simulation is a valuable tool which allows for experimentation and validation within a competitive airline environment. As Gorin and Belobaba[73] describe, analytic RM models entail a certain level of simplification due to their static nature, oftentimes leaving such models ill-prepared for competitive actions among airlines or arbitrary passenger booking behaviors. By taking a simulation approach, a dynamic representation of RM practices can be modeled in a competitive framework characterized by realistic interactions between passengers' booking decisions and RM systems.

This chapter contains an overview of the Passenger Origin-Destination Simulator (PODS) used to test Hybrid Forecasting (HF) for this thesis. Here we describe the PODS component modules, including passenger choice, forecasting, and seat inventory control methods; we also present the simulated air transportation network used for experimentation, consisting of the competing OD services and fare structures offered by two airlines. Finally, we describe the specific experiments to be performed with HF in combination with Fare Adjustment (FA) and Path Categorization (PCAT).

3.1 PODS Background

3.1.1 Introduction

Originally developed by C. Hopperstad, M. Berge, and S. Filipowski at the Boeing Company, PODS is an evolution of Boeing's earlier Decision Window Model (DWM)[22] for passenger choice, and is a computer simulator of competitive airline networks. In the 1990s, the Massachusetts Institute of Technology and several major international airlines created the PODS Research Consortium – a partnership which employs the tool to study, develop, and test new RM techniques.

Unlike other RM simulators, PODS does not model passenger demand as a computationally simple independent variable (analogous to Lee's micro level as described in Section 2.1), but rather as the more realistic aggregation of millions of passenger level choices among competing airlines, schedules, and fare products. As computing power has increased, so has the efficacy of PODS to simulate airline competition over larger, more complex networks using a greater repertoire of RM techniques.

At its most basic, PODS can be described as a simulation of the interactions of two groups in a user-defined transportation network: airlines and passengers. The airline side of the simulator consists of a third generation RM system, similar to that described in Section 2.2, which provides air travel offerings to consumers. And on the other side of the fence are simulated passengers seeking air travel in a specific

[73] Gorin, T., P. Belobaba. 2004. Revenue management performance in a low-fare airline environment: insights from the Passenger Origin-Destination Simulator. *Journal of Revenue and Pricing Management.* Volume 3, Issue 3, pp. 215-236.

OD market and trying to decide among multiple airlines, paths, and fare classes. This dichotomy is shown below in Figure 9.

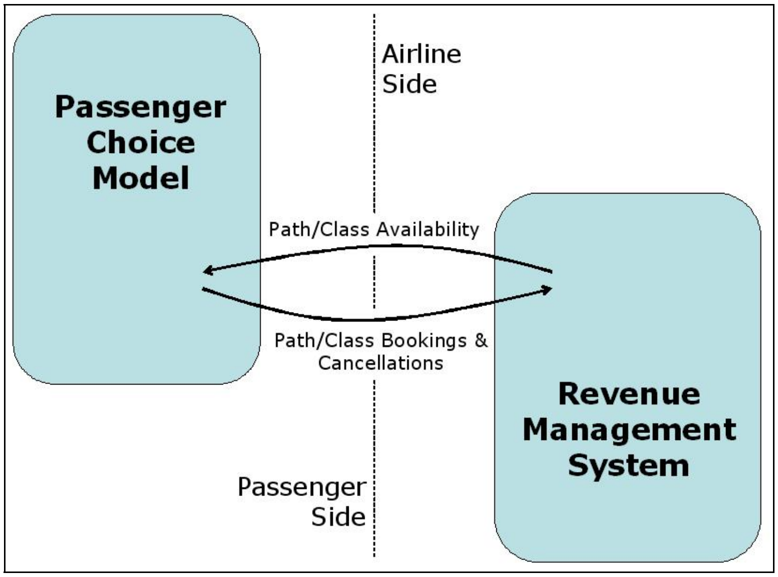

Figure 9: PODS Structure

From the customer standpoint (or the demand side), each prospective passenger decides upon some combination of carrier, path, and fare class (or does not purchase) to fit his or her needs; this booking information is passed over the fence to the airlines, which collect the revenue and record the bookings for future analysis. The PODS passenger choice model is described further below in Section 3.1.2.

From the airline perspective (or the supply side), the RM system uses the current booking levels as well as historical booking patterns to determine the availability of fare classes on each path; this availability information is passed over the fence to the passengers in an effort attract future bookings. The traditional RM components within PODS are described further in Section 3.1.3 below; the newer developments within PODS for less restricted fare structures are described further in Section 3.2.

In PODS, the booking process stretches over 16 successive time frames, the first beginning 63 days prior to departure, and the sixteenth ending on the departure date. Passenger events such as bookings and cancellations are randomly spread within each of these time frames, while the RM system's major inventory actions typically occur at the start of each one (though certain PODS features modify availability throughout the entire booking process, regardless of time frame). In the experiments for this thesis, the time frames initially last for one week, but shrink to two days as departure nears in order to capture the expected increase in booking activity, as shown below in Table 2.

Table 2: User-Defined Time Frames

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

Each PODS simulation actually consists of thousands of simulations averaged together to provide overall operating statistics for each simulated airline on a per day basis. In PODS terminology, a "run" consists of "trials" and "samples." For the experiments in this thesis, a run – or an individual simulation – is the average of five independent trials; and each trial is the iterative result of 600 samples. Each sample represents a single departure day, and the reason for the large sample size is to ensure statistical significance of a simulation's results.

Because the starting values for each of the five trials (i.e. sample 1 of 600) are arbitrarily user defined, and each sample has some degree of correlation to the previous one as described below in Section 3.1.3.1, we discard the first 200 samples as the user inputs are gradually replaced with calculated values from simulation. So the results of each 600 sample trial are based only on the last 400 samples, and every PODS run is actually the averaged result of 2,000 daily simulations.

3.1.2 PODS Passenger Choice Model

In the airline industry, the success of a given RM technique depends directly upon passenger response (or non-response) to that technique. And in the PODS simulator, those responses are governed by a particular choice model, as briefly described in this section. A comprehensive discussion of the PODS Passenger Choice Model, including its assumptions, logic, and ultimate validation is provided by Carrier[74].

As previously mentioned, PODS itself, and especially its Passenger Choice Model, is an outgrowth of Boeing's original Decision Window Model[22], but with several notable enhancements. Included among these changes to the original DWM is passenger consideration of fares on each path and recognition of advance purchase (AP) restrictions assigned to certain fare classes[32]. The PODS Passenger Choice Model can be divided into four sequential steps: Demand Generation, assignment of Passenger Characteristics, definition of a Passenger Choice Set, and a specific Passenger Decision. This structure is depicted below in Figure 10; implicit in this framework are the links to the RM system, as shown above in Figure 9.

[74] Carrier, E. 2003. Modeling airline passenger choice: passenger preference for schedule in the Passenger Origin-Destination Simulator (PODS). Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

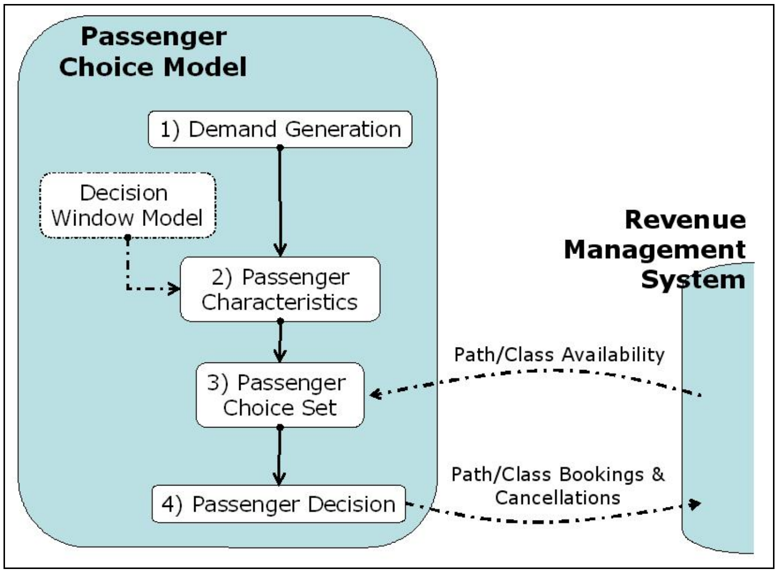

Figure 10: Passenger Choice Model in PODS

3.1.2.1 Demand Generation

In this initial step, the average daily air travel demand is generated for each OD market in the user-defined network. While this estimate was formerly rooted in a gravity model between the origin and destination cities, the OD demand is now based upon data provided by the PODS Consortium's airline members. For the experiments in this thesis, the total passenger demand is apportioned between leisure and business passengers, 65% and 35%, respectively, according to recent US airline industry data. Next, the Passenger Choice Model randomly generates variability around this average daily demand; however, the demand generation process does not include seasonal or day-of-week variability. Finally, the passenger arrival patterns through the booking process are modeled for both business and leisure segments according to user-defined booking curves; the booking curves used for experiments in this thesis are shown below in Figure 11.

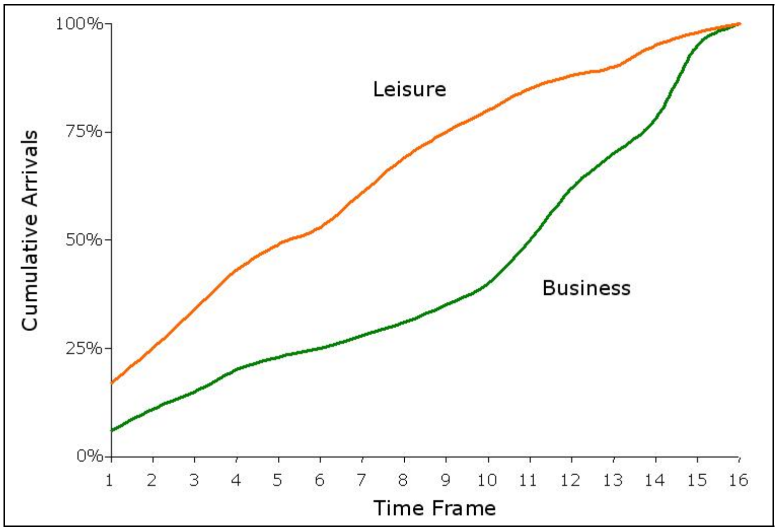

Figure 11: PODS Booking Curves

In PODS, the intensity of bookings is governed by a "demand multiplier" (DM), which can be used to simulate periods of low and high demand. In this thesis, the default DM used in experiments was 0.9, with a value of 0.8 used for low demand situations, and a value of 1.0 used for high demand situations. The particular experiments in which the DM varied are described in greater detail in Section 4.2.

3.1.2.2 Passenger Characteristics

In this second step, three specific characteristics are assigned to the individual passengers generated in Section 3.1.2.1: a decision window, a maximum willingness-to-pay (WTP), and a set of disutilities associated with certain aspects of his or her booking, such as the potential fare restrictions.

1. The decision window for each passenger is defined by that passenger's earliest acceptable departure time and latest acceptable arrival time; business travelers tend to have smaller decision windows than their leisure counterparts reflecting their time sensitivity. At this stage of the booking process, all path and fare class combinations which are feasible within a passenger's decision window are considered equally appealing, and all infeasible paths are equally unappealing and require re-planning of the decision window.

2. As its name implies, each passenger's maximum WTP reflects the maximum out of pocket fare that he or she willing to pay for OD travel. Any fare exceeding the WTP will be excluded from his or her consideration. As a result, any fare above the maximum WTP of the passenger will be excluded from his choice set. These WTP values are taken from a user-defined price-demand curve, examples of which are shown below in Figure 12 for leisure and business passengers (the business passenger curve is typically flatter, representing less price sensitivity).

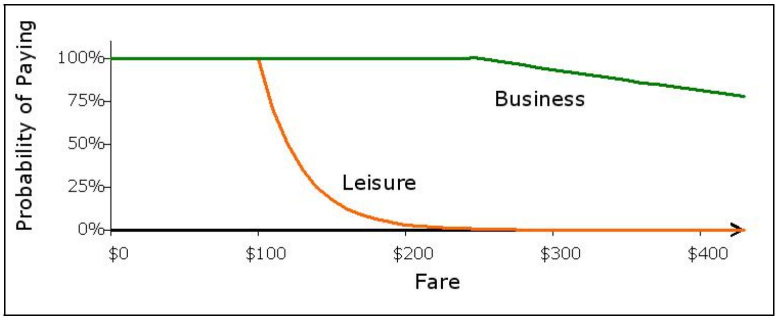

Figure 12: Sample Passenger WTP Curves in PODS

3. A set of disutility values is assigned to each passenger, which represent his or her sensitivity to fare product restrictions, schedule preference (re-planning disutility, as mentioned above), and path quality (non-stop versus connecting paths). In terms of booking restrictions, passenger disutility distributions are typically defined for Saturday night stay requirements, itinerary change fees, and non-refundability constraints. For a generated passenger within PODS, his or her disutilities are randomly selected based upon these user defined passenger type probability distributions. Lee[75] provides more detailed descriptions of the passenger disutility assignment process in PODS.

3.1.2.3 Passenger Choice Set

In the next step of the PODS Passenger Choice Model, each passenger is presented with a set of fare products to consider. The maximum size of this choice set is determined by the number of airlines, the network size, and the number of fare products offered per airline. Some of the alternatives will immediately be eliminated from the passenger's choice set for the following reasons:

- The RM system for one (or more) airline has closed a particular set of fare classes and/or paths in the desired OD market. This requires acquiring availability data from the airline side of the PODS simulator.
- Advance purchase requirements are not met for certain fare products.
- A particular fare exceeds the passenger's WTP.

It is also important to realize that the "do-nothing" option of not booking is always an option available to each passenger.

3.1.2.4 Passenger Decision

Faced with this choice set, the passenger must make a decision to either purchase a particular (available) fare product, or not to book. The "total generalized cost" of each available alternative (including the do-nothing) is calculated by summing the fare and the relevant disutilities described in Section 3.1.2.2, and each passenger selects the option with the lowest generalized cost. At this point, the Passenger-Airline fence in Figure 9 is crossed again, as the booking decision is returned to the

[75] Lee, S. 2000. Modeling passenger disutilities in airline revenue management simulation. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

respective airline's RM system. The airline decreases its available seat inventory and records the booking details in its historical database to aid in future demand forecasts.

3.1.3 RM System within PODS – Traditional Components

On the airline side of the PODS program, each user-defined carrier is represented by a third generation RM system similar in structure to the generic arrangement previously discussed (and illustrated in Figure 4). In an effort to better isolate the effects of forecasting, seat allocation optimizers, and other RM techniques, an overbooking component has been excluded from the simulator. The PODS RM system consists of three interacting components: the Historical Booking Database, the Forecaster, and the Seat Allocation Optimizer. This structure, as well as the relationships among components, is depicted below in Figure 13; implicit in this framework are the links to the PODS Passenger Choice model, as shown above in Figure 9.

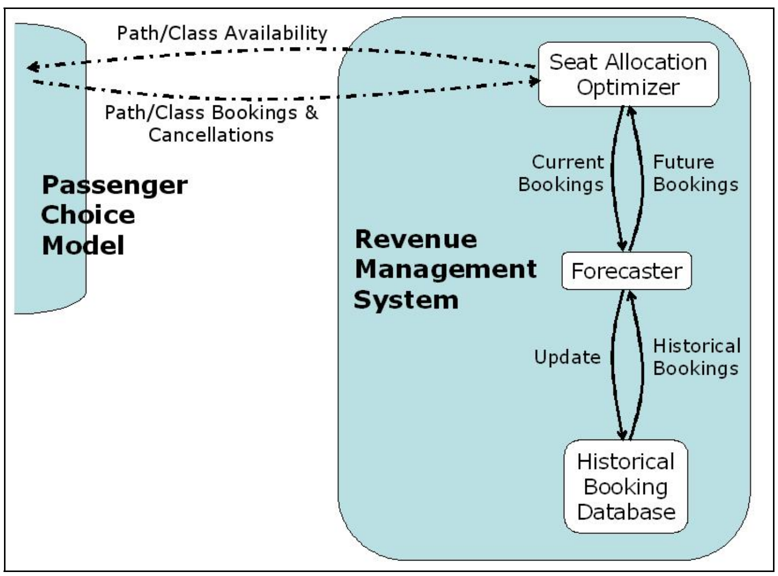

Figure 13: PODS RMS Structure

3.1.3.1 Historical Booking Database

Analogous to the top level of Figure 4, the PODS Historical Booking Database is a repository for each respective carrier's booking data by fare class and path. It is this database that is initially filled with user defined default values at the onset of each trial; as bookings occur, the actual data gradually replace the initial default values. Thus, the need exists to discard the first 200 samples of each trial, as the replacement process enables us to minimize the correlation effects of these initial values absent booking data. For the experiments performed in this thesis, the Historical Booking Database is limited to the previous 26 previous simulated departures of a particular service.

3.1.3.2 Forecaster

One of the three Modeling Components of Section 2.2's generalized RM system, the forecasting mechanism within PODS extracts data directly from the Historical Booking Database in an effort to estimate future demand for a given path and fare class. The PODS user typically chooses one of the two forecasting methods: regression or pick-up

As mentioned above in Sections 2.2.1.1 and 2.5.1, classical pick-up forecasting is used heavily in this thesis as the product-oriented component of Hybrid Forecasting. A pick-up model of demand generates a forecast of total bookings by adding the average of the incremental historical bookings to whatever current bookings have already been taken at a certain point prior to departure; that is, the final number of bookings depends on the number of current bookings in hand and the average number of bookings "picked up" between the current reading day and departure. This relationship is expressed below.

$$B_{F,i} = \left[\frac{1}{f} \sum_{j=1}^{f} w_j \left(B_{j,i} - B_{j,i-n} \right) \right] + B_{F,i-n}$$

Where:

- $B_{F,i}$ signifies the total number of bookings after time frame i for flight F;
- F signifies the flight for which the demand forecast is desired;
- n signifies the number of time frames over which the pick-up is calculated;
- f signifies the number of flights upon which the demand forecast is based;
- j signifies the index of flights upon which the demand forecast is based;
- w_j signifies is the weighting value applied to any particular flight (optional – not used in this thesis);

But forecasting upon the raw data is typically not sufficient for estimating demand in an unbiased fashion. Data taken from historical bookings are often referred to as "censored" or "constrained," because they reflect only actual bookings made while a fare class showed seat availability – not the hypothetical number of bookings that would have materialized had that fare class remained open indefinitely (again, assuming demand independence among fare classes).

To overcome this bias towards underestimating demand, one of several "detruncation" techniques can be applied to estimate unconstrained historical demand. In PODS, the two most common techniques are Booking Curve and Projection detruncation. As illustrated in Figure 14, Booking Curve detruncation uses a percentage-based multiplier which extrapolates demand for closed fare classes using trend data from open ones. Booking Curve detruncation was used for all experiments in this thesis.

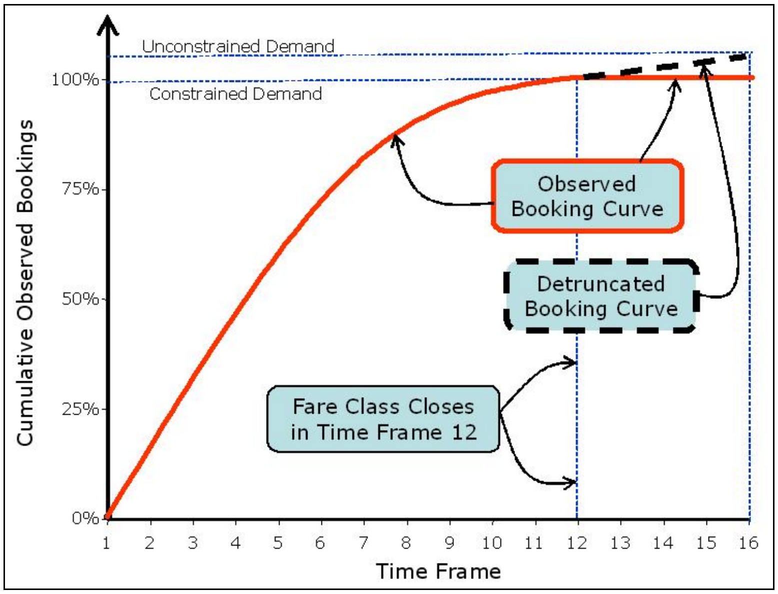

Figure 14: Booking Curve Detruncation Example

More thorough discussions of the various forecasting and detruncation methodologies in PODS (including experiments performed with their variants and combinations) is provided by Zickus[29], Skwarek[30], Usman[31], Gorin[32], and Wickham[33], among others.

3.1.3.3 Seat Allocation Optimizer

In PODS, the user defines the seating capacity for each leg within the network (100 seats on all legs for all experiments in this thesis). And the Seat Allocation Optimizer determines the manner in which that fixed inventory of seats is made available to passengers requesting travel. In the PODS RM system, this optimizer can take a variety of forms with various levels of sophistication.

At its most basic level – the First Come First Served (FCFS) method – the inventory optimizer actually performs no optimization, and simply allows passengers to book seats in any fare class they request on a space available basis. This method is typically used only for certain baseline comparisons, or to simulate the actions of an irrational carrier using a "dumb" RM system. Slightly more refined, PODS also has a class of threshold algorithms which can limit the number of bookings in a specific class to a Fixed Threshold (FT) ratio of overall bookings, or even aim for a target load factor using an Adaptive Threshold (AT) methodology.

The PODS Fare Class Yield Management (FCYM) scheme uses variants of the EMSR heuristic to determine the optimal protection levels for each leg in the system, as described above in Section 2.2.3.1. Under the assumption of independent demand by fare class, EMSRb determines the nested booking limits by protecting incremental seats in each class as long as their EMSR exceeds that of the class beneath it. So for an example flight leg with four fare classes and a 95 seat capacity shown in Table 3, Y class is protected as long its EMSR ($525 multiplied by the probability of selling just one more Y seat) exceeds the EMSR of B class ($400 multiplied by the probability of selling one B seat). Of course, these calculations are straight forward when demand is assumed to follow a Gaussian distribution; and the mean and standard deviation values for each fare class and leg are gathered from the PODS forecaster, as previously described in Section 3.1.3.2.

Table 3: FCYM Example – Booking Capacity of 95

Class	Fare	Demand μ	Demand σ	Booking Limit
Y	$525	18	5	95
B	$400	25	8	81
M	$305	34	10	56
Q	$195	38	12	18

For this example, the point of indifference occurs between the 14th and 15th Y class seat; by repeating this procedure for each fare class until the leg's booking capacity is reached, the optimal booking limits can be determined as shown in the EMSR curves of Figure 15.

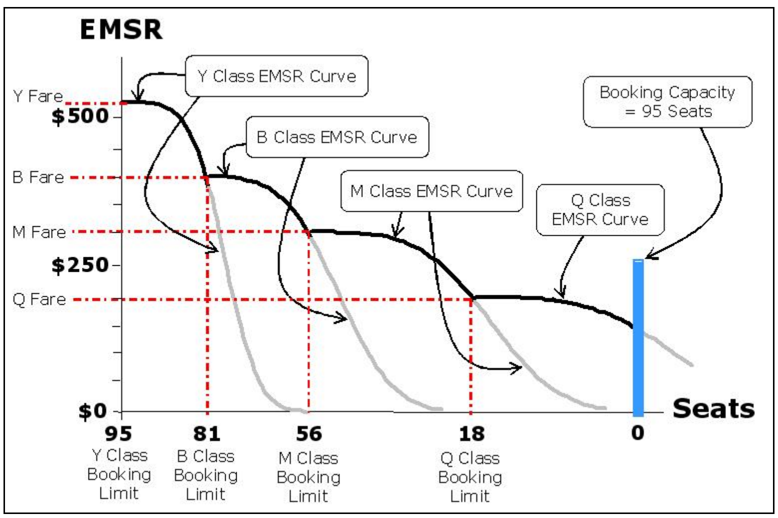

Figure 15: FCYM Example EMSR Curves and Associated Booking Limits

The booking limit for each class is the maximum number of seats the airline should sell for that particular class. In the given example, if 16 Q class seats have already been sold, the airline will not sell more than 2 additional Q seats when using EMSRb; and if the limit of 56 total M and Q class seats has been reached, the airline will only offer Y and B class seats for the remainder of the booking period. The entire nesting structure – including the protections for each fare class – for this particular example is shown in Figure 16.

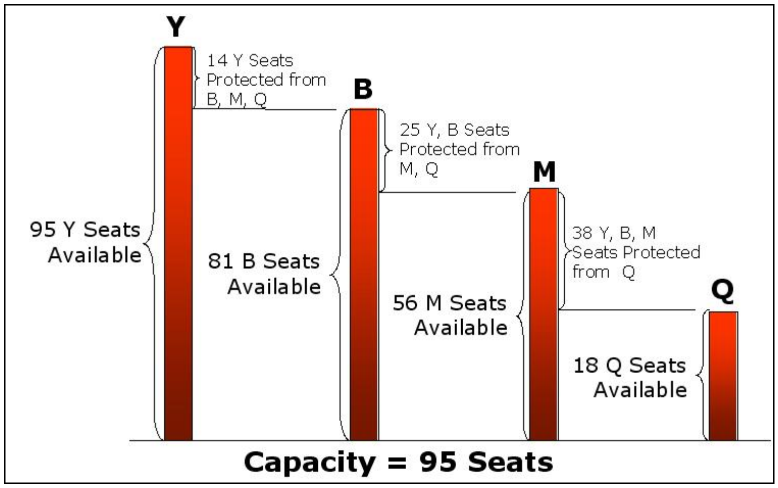

Figure 16: FCYM Example Nesting Structure and Associated Booking Limits

In terms of OD seat allocation, one strategy available in PODS involves the virtual nesting of all local and connecting fare classes sharing a particular leg into a set of buckets. Each bucket can contain one or more fare classes, with the user determining the bucketing structure for the simulation. Within PODS, the simplest implementation is known as Greedy Virtual Nesting (GVN) in which the connecting and local fare classes are sorted on the basis of overall fare, and the EMSRb logic is used to generate protections for each virtual nest on each leg throughout the network.

A simple example of GVN for one 95 seat leg in a two leg network (so 2 local paths, 1 connecting path) with eight virtual nests is shown below in Table 4. This particular example is small enough that each fare class is assigned its own bucket. And in the same manner as the FCYM example above, the optimal booking limits for this leg can be determined as shown in the EMSR curves of Figure 17.

Table 4: GVN Example – Booking Capacity of 95

	Class	Fare	Virtual Nest	Demand μ	Demand σ		Virtual Nest	Booking Limit
Local	Y	$525	V2	10	3		V1	95
	B	$400	V5	12	5		V2	88
	M	$305	V6	18	5		V3	78
	Q	$195	V8	16	6		V4	69
Connecting	Y	$555	V1	12	3		V5	59
	B	$480	V3	10	5		V6	43
	M	$435	V4	12	6		V7	23
	Q	$250	V7	12	9		V8	8

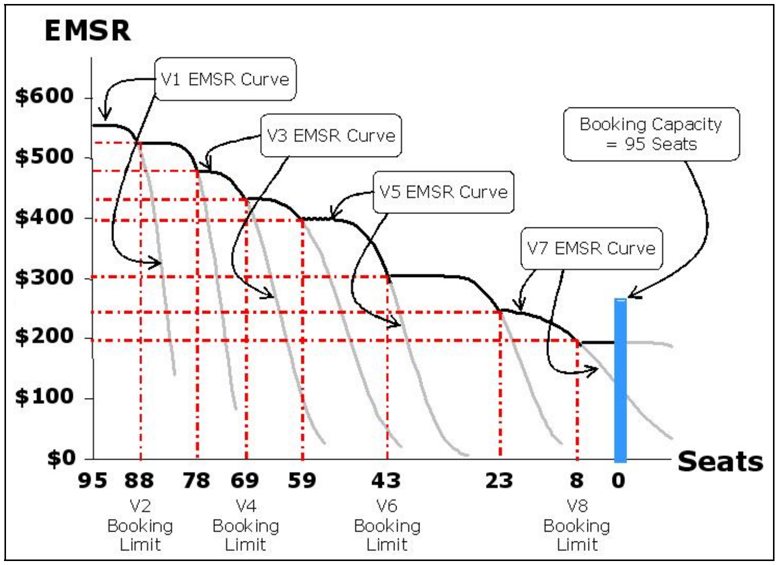

Figure 17: GVN Example EMSR Curves and Associated Booking Limits

Under GVN, an airline will accept a booking only if the requested fare class shows availability on each of its legs. So for the above example, if Virtual Nest 4 has reached its booking limit of 69 seats, the carrier will not offer any M class connecting service, even if the second leg for this connecting path has no bookings whatsoever.

And herein lies the limitation of GVN: being too greedy. By blindly protecting the fare classes on each leg which produce the highest expected revenue, regardless of local or connecting service, GVN ignores the opportunity cost of a connecting passenger, or the value of a local passenger being displaced on one leg due to accepting a connecting passenger on a different leg.

As previously mentioned in Section 2.2.3.2, this thesis makes heavy use of the Displacement Adjusted Virtual Nesting (**DAVN**) method of seat allocation, which accounts for displacements within the network by way of applying a revenue penalty, or a displacement cost, to connecting fares. The first step in calculating these displacement costs is the generation of demand forecasts. And in contrast to the inventory optimization schemes described above, DAVN is one of the few methods to utilize path-based demand forecasting for each fare class.

DAVN then solves a deterministic linear program (LP) to determine the optimal number of passengers for each OD path and fare class; the formulation of this LP, as explained by Williamson[43], is shown below.

$$Maximize \sum_{OD} \sum_{f} fare_f^{OD} x_f^{OD}$$

Subject to two constraints:

1. $x_f^{OD} \leq \mu_f^{OD}$ $\quad\quad$ $\forall$ fare classes, *f*, $\forall$ paths, *OD*;

2. $\sum_{OD}\sum_{f} x_f^{OD}\delta_{OD}^{j} \le C_j \quad \forall$ fare classes, f, $\forall$ flight legs, j;

Where:

- OD signifies a particular path, from Origin to Destination;
- f signifies a fare class on a path;
- j signifies a flight leg;
- x_f^{OD} is the decision variable; signifies the optimum number of passengers in the fare class, f, on the path, OD;
- $fare_f^{OD}$ signifies the revenue the carrier collects from a passenger booking in the fare class, f, on the path, OD;
- μ_f^{OD} signifies the mean forecasted passenger demand for the fare class, f, on the path, OD;
- C_j signifies the maximum number of passengers, or the capacity, for the flight leg, j; this value is 100 for all experiments in this thesis;
- δ_{OD}^{j} is an indicator variable; equals 1 if the leg, j, is on the path, OD; 0 otherwise;

The objective function maximizes the revenue (the fare each passenger paid multiplied by the number of passengers, for all fare classes in all OD paths), subject to two constraints. The first constraint limits the number of passengers booking in a particular fare class to the mean demand forecast.

The second constraint limits the number of passengers on an aircraft (or the sum of the passengers in each fare class on each path sharing that particular flight leg). This constraint is vital to the DAVN algorithm because the LP's dual solution reveals the value of increasing C_j by one seat for each leg, j, in terms of additional revenue to the airline. This value is denoted as d_j.

Known as the "shadow price" for each leg in the network, d_j can be interpreted as the minimum value the carrier would be willing to accept for an additional seat on leg j, or the expected network revenue increase when relaxing the capacity constraint on the leg. In PODS, this shadow price is also the displacement cost used to penalize connecting fares for preventing the booking of local passengers elsewhere in the network.

The next step in the DAVN process is the calculation of "pseudo fares" for each fare class f on leg x in network path OD by subtracting the displacement cost from each fare, as shown below.

$$Pseudo\ Fare_f^{x,OD} = fare_f^{OD} - \sum_{\substack{j \in OD \\ x \neq j}} d_j$$

For a particular itinerary through the network, OD, the pseudo fare for a given leg on that itinerary is the path's regular fare minus the shadow price for all other legs on that itinerary. So for local passengers, the pseudo fare is equal to the actual fare to be paid; for connecting passengers, the pseudo fare is the difference between the actual fare on the leg the passenger is booking and the shadow price of the other leg (in a two-leg case).

The final step in DAVN is to bucket all the fare classes (both local and connecting) sharing each leg in terms of pseudo fares, and to determine the seat protection limits in terms of these buckets just as in GVN. Thus, despite forecasting on the OD level, seat allocations in DAVN are still performed on the leg level. But unlike in GVN, the bucketing of fare classes by pseudo fares (not regular fares) provides a correction for passenger displacement throughout the network.

Beyond virtual nesting, PODS also has a class of OD optimizers that take a bid price (BP) approach, and several that utilize dynamic programming (DP). For more detail on BP methods within PODS, including NetBP, ProBP, and HBP, see the sources described in the literature review, Section 2.2.3.2. And Vanhaverbeke[76] discusses the DP approach to RM and its use within PODS in this Master's thesis. Neither BP nor DP based optimizers were used in the experiments for this thesis.

3.2 RM Developments within PODS for Simplified Fare Structures

As discussed the previous section, PODS is a complex and valuable tool created to study RM issues in a competitive airline environment. Yet, PODS was developed in a different competitive environment, before the pervasive growth of LCCs, as described above in Chapters 0 and 0. In less restricted fare structure environments, carriers cannot rely on product restrictions to fence demand into neatly defined fare classes anymore. In an effort to encourage more sell-up to higher classes, several different methods have been implemented and tested within PODS.

3.2.1 Traditional Sell-up in PODS

Belobaba and Weatherford[77], in recognition of the revenue gains to be had by "selling up" passengers from lower fare classes to more expensive products (and in tacit acknowledgement of the weakness of the fare class demand independence assumption) developed a sell-up algorithm for use with EMSRb (and thus, GVN and DAVN when applied to virtual buckets on a leg). Implemented in PODS and extended to the various RM techniques available, the Belobaba-Weatherford heuristic strengthens the Seat Allocation Optimizer's protections of higher classes in order to account for passengers selling up to these classes.

Designed to account for sell-up in unrestricted environments, this technique and its use within PODS is described in greater detail by both Gorin[32] and Skwarek[30]. It is important to note that the Belobaba-Weatherford heuristic does not actually estimate the WTP within a fare class or virtual nest. Rather, PODS requires a user input sell-up probability by time frame. Thus, the effectiveness of the sell-up protections is highly dependent on the estimate of sell-up utilized by a specific airline (i.e. input by the user) – a recurring challenge in PODS.

---

[76] Vanhaverbeke, T. 2006. Revenue management based on dynamic programming in unrestricted and simplified fare structures. Master's thesis. Massachusetts Institute of Technology, Cambridge, MA.
[77] Belobaba, P. P., L. R. Weatherford. 1996. Comparing decision rules that incorporate customer diversion in perishable asset revenue management situations. *Decision Sciences*. Volume 27, Issue 2, pp. 343-363.

3.2.2 Q-forecasting

PODS utilizes the technique known as "Q-forecasting," described in Section 2.4.1, as a method of managing passenger sell-up and counteracting the spiral-down effect in the restriction-free environments where all passengers will book in the lowest priced fare class. In short, Q-forecasting entails using historical booking data to estimate the number of potential future arrivals in this lowest priced fare class, then converting that value into an equivalent number of potential bookings in each of the higher fare classes. Thus, the optimizer strategically limits seat availability by fare class to induce sell-up from lower to higher classes and drive the carrier's revenue upward based on passenger WTP. But as with the Belobaba-Weatherford sell-up heuristic described above in Section 3.2.1, the Q-forecasting methodology does not actually estimate WTP, but requires a user input estimate of sell-up. For Q-forecasting, this sell-up input value is known as a "**FRAT5**."

3.2.2.1 Use of "FRAT5" Values

Within the PODS Q-forecaster, sell-up is governed by a FRAT5 value, or the fare ratio between a low and high fare class which entices 50% of the demand for the lower class to sell-up to the higher class. More specifically, the FRAT5 is a proxy for passenger WTP which quantifies the probability a passenger will sell-up from Q-class to some more expensive fare class. Thus, a low FRAT5 value denotes high price sensitivity among passengers, and vice versa. In PODS, sell-up is assumed to follow an inverse exponential shape, as shown below in Figure 18. This shape has been chosen based upon empirical observation, intuitive expectation, and ease of computation.

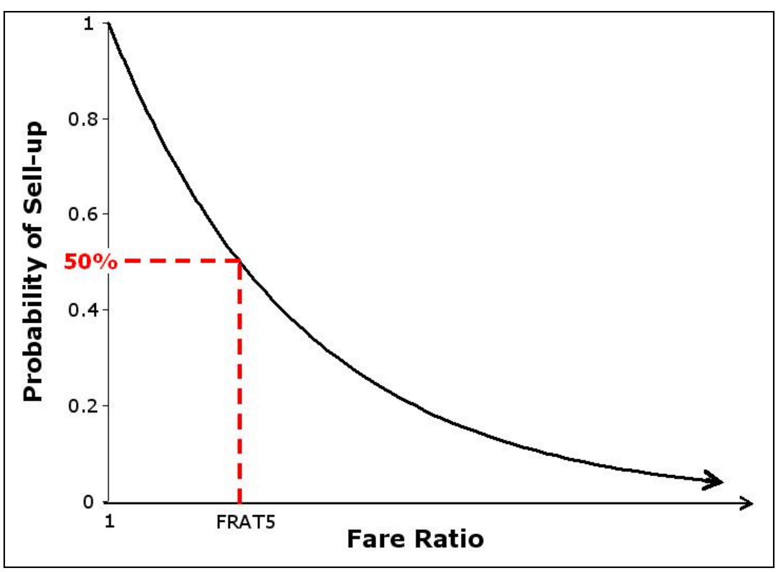

Figure 18: Inverse Exponential Form of Sell-up Probability versus Fare Ratio

So the probability of sell-up, $psup$, from Q-class to some higher fare class, f, is an inverse exponential function of the fare ratio between Q and f, and a sell-up constant, $supcon$, based on the FRAT5, as shown below.

$$psup_{Q \to f}\left(fare_f\right) = \exp\left(-supcon \cdot \left(\frac{fare_f}{fare_Q} - 1\right)\right)$$

Where:

- $supcon = -\frac{\ln(0.5)}{FRAT5 - 1}$
- $fare_f$ signifies the fare of the higher fare class, f;
- $fare_Q$ signifies the fare of the lowest fare class, Q;
- $FRAT5$ signifies the fare ratio at which 50% of passengers will sell-up from fare class Q;

As the FRAT5 values increase, the probability of sell-up to higher fare classes increases, as shown below in Figure 19. Thus, passenger behavior becomes more aggressive with higher FRAT5s, and a simulated airline in PODS using a high FRAT5 value will assume that passengers will demonstrate high WTP and will protect more high fare class seats to account for expected sell-up.

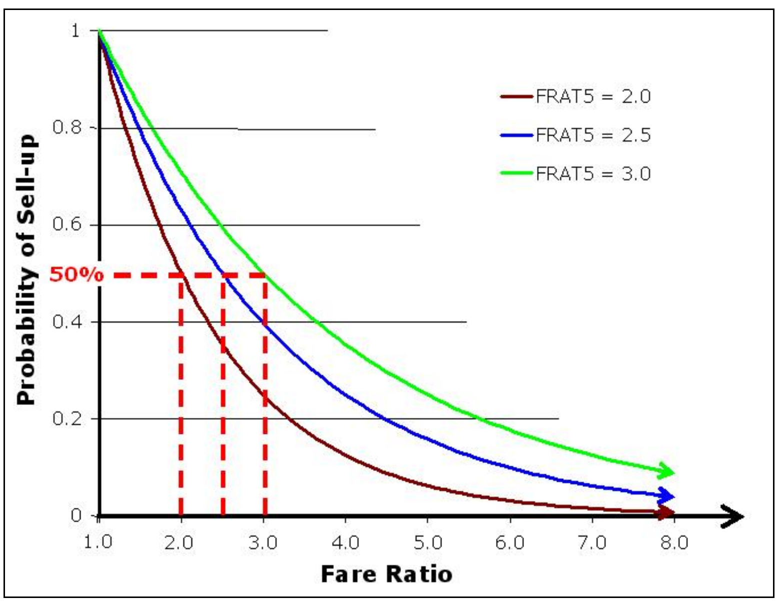

Figure 19: Probability of Sell-up by FRAT5 Value

But just as we assume that sell-up increases with higher FRAT5 values, we also assume that passenger sell-up is more prevalent closer to departure, as WTP increases (i.e. price sensitivity decreases throughout the booking process). In PODS, we capture this behavior by assuming that FRAT5 values gradually increase as departure draws near. We also assume that this increase generally follows an "S-shape," as shown below in Figure 20, where the rate of FRAT5 increase is higher between time frames 8 and 11 than it is at the beginning and end of the booking process.

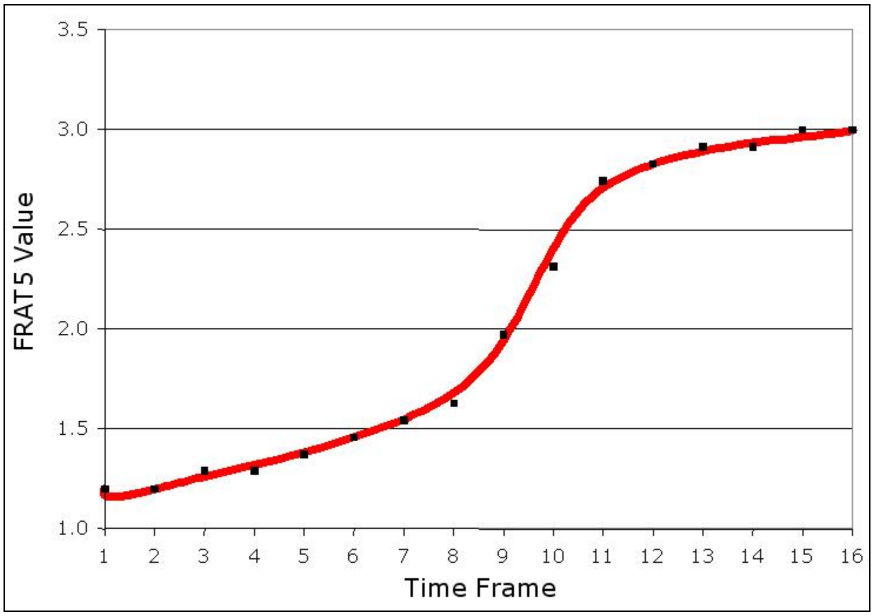

Figure 20: S-Shape of FRAT5 (for FRAT5 Series "C")

The use of FRAT5s represents the WTP assumptions airlines must make for their internal RM systems – not a change to the PODS Passenger Choice Model. The underlying WTP for simulated passengers does not change.

For simulations of alternative scenarios in PODS, we use 17 distinct FRAT5 series, each assuming a different level of price sensitivity among passengers. For each simulated airline using Q-forecasting, the PODS user must select a specific FRAT5 series which the airline will assume represents passenger behavior. These 17 different series are labeled ("A9" through "A1", and then "A through H"), and each takes the characteristic "S-shape" as shown in Figure 20 (which depicts FRAT5 series "C"). The FRAT5 values for Time Frame 1 and Time Frame 16 for each of the 17 series are shown below in Table 5.

Table 5: FRAT5 Series

FRAT5 Series	A9	A8	A7	A6	A5	A4	A3	A2	A1	A	B	C	D	E	F	G	H
Time Frame 1	2.5	2.4	2.3	2.2	2.1	2.0	1.9	1.8	1.7	1.6	1.4	1.2	1.15	1.1	1.08	1.05	1.02
Time Frame 16	8.5	8.0	7.5	7.0	6.5	6.0	5.5	5.0	4.5	4.0	3.5	3.0	2.5	2.0	1.8	1.5	1.2

Clearly, A9 represents the most aggressive passenger behavior (highest FRAT5s), and H represents the least aggressive sell-up due to its lowest FRAT5s. The FRAT5 values by time frame for six arbitrarily selected series (A4, A2, A, C, D, F) are shown below in Figure 21.

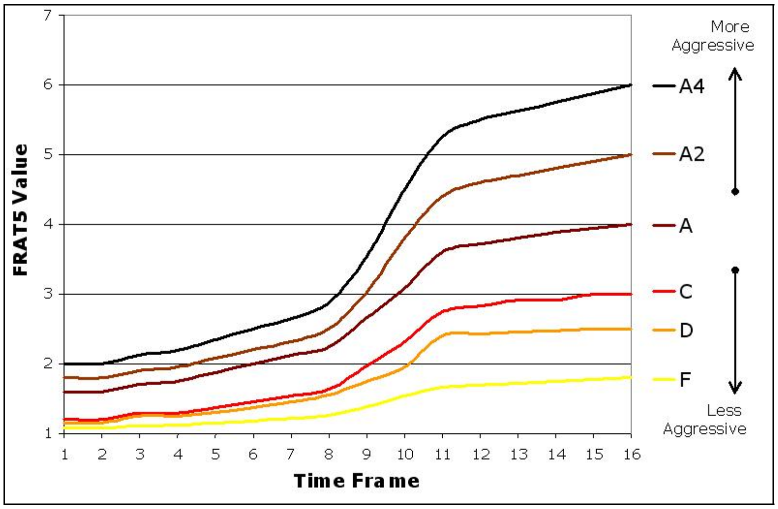

Figure 21: Selected PODS FRAT5 Values by Time Frame

Combining these time series FRAT5 values from Figure 20, and the sell-up probabilities by fare ratio (FRAT) from Figure 19, we can observe the probability of sell-up in each time frame at different fare ratios, as shown in Figure 22 (an example using FRAT5 series "C"). This particular example also demonstrates the sell-up probabilities from Q-class to three higher classes, with fare ratios of 2.0, 3.0, and 4.0 to $fare_Q$. As shown below, sell-up is always less for higher fare classes, as we would expect in an absence of fare class restrictions; and the probability of sell-up increases closer to the departure date, as previously mentioned.

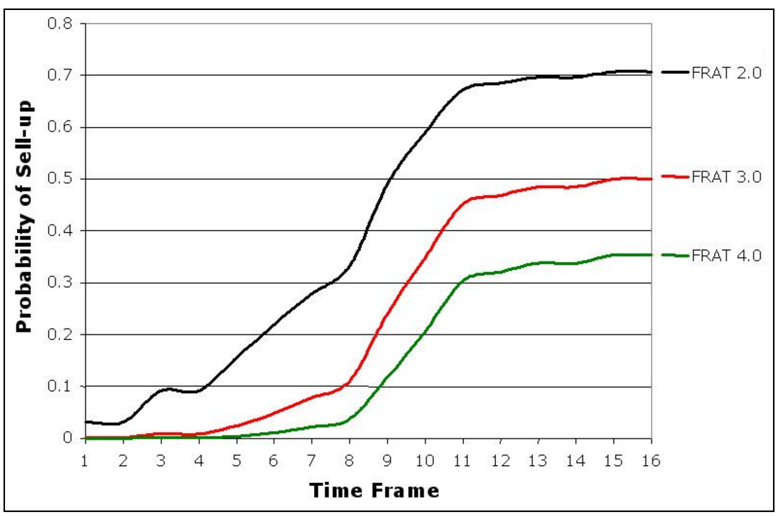

Figure 22: Probability of Sell-up in Time Frame by Fare Ratio (Assuming FRAT5 Series "C")

3.2.2.2 Q-Forecasting Methodology in PODS

As previously mentioned, Q-forecasting was designed for use in the restriction-free environments characterized by multiple fare classes identical except for price; it is in these environments where airlines are vulnerable to spiral-down, as previously mentioned in Section 2.3.2. As with the traditional forecasters described in Section 3.1.3.2, Q-forecasting relies on detruncated historical data. And in unrestricted fare structures, booking data in each time frame can only exist for a single fare class because it is impossible for demand to materialize in higher fare classes with no booking restrictions to segregate passengers. However, the use of observations from multiple departures allows the forecaster to gather data points in several fare classes in the same time frame.

We illustrate Q-forecasting using a simple example with four unrestricted fare classes, priced at $400, $300, $200, and $100, as shown below in Figure 23. In this example, we are forecasting for Time Frame *i*, and we have historical booking demand for Time Frames *i+1*, *i+2*, and *i+3*. For example, in Time Frame *i+1*, the first observation of unconstrained demand in fare class 2 is 3 bookings, based upon historical data.

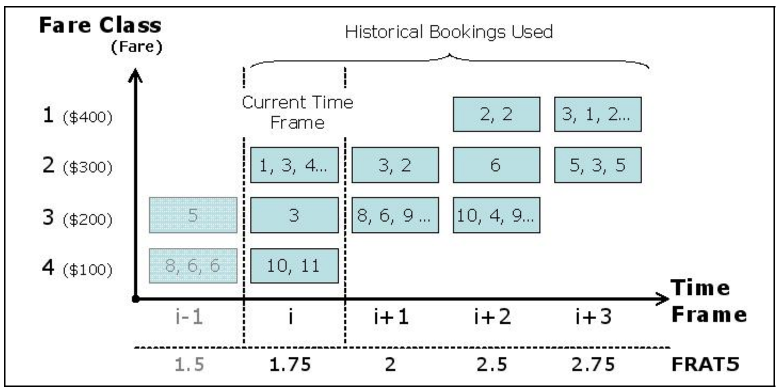

Figure 23: Example Historical Bookings in Unrestricted Fare Structure Environment

The first step in Q-forecasting is to use FRAT5s – as described above in Section 3.2.2.1 – to determine the probability of sell-up from the lowest class for each fare class and each time frame. This stage is illustrated below in Figure 24 for our basic Q-forecasting example. Note that this process is independent of the actual historical bookings, and that these probabilities depend entirely on FRAT5 values for each time frame and the fare of each class – a dependence relationship depicted by the arrows in Figure 24.

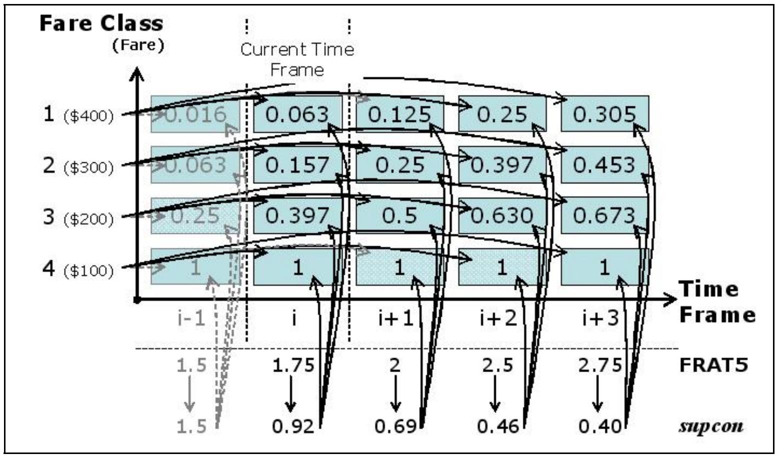

Figure 24: Example Calculation of Sell-up Probability from Fare Class 4 to all other Fare Classes by Time Frame

Using these sell-up probabilities for each time frame and fare class, the next step in Q-forecasting is calculating the equivalent number of Q-bookings for the historical booking observations in each fare class using the formula below.

$$hbk_{Q \to f,tf} = \frac{hbk_{f,tf}}{psup_{Q \to f,tf}}$$

Where:

- $hbk_{Q \to f,tf}$ signifies the estimated equivalent demand for fare class Q in fare class f in time frame tf;
- $hbk_{f,tf}$ signifies the mean unconstrained demand in fare class f and time frame tf;
- $psup_{Q \to f,tf}$ signifies the probability of sell-up from fare class Q to f in time frame tf;

For our simple example, the equivalent Q-bookings in each fare class and time frame are shown below in Figure 25, as are the total number of estimated Q-bookings in each time frame.

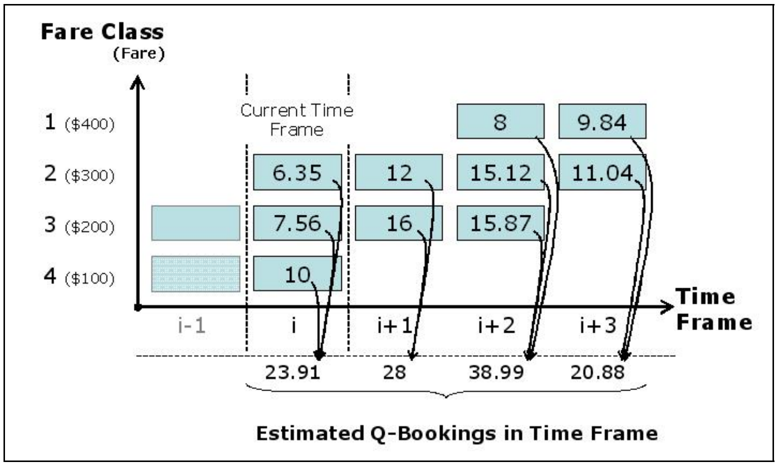

Figure 25: Example Equivalent Q-Bookings in Time Frame

The next step is forecasting total Q-Bookings in each future time frame by applying traditional forecasting techniques to the series of estimated Q-bookings in each time frame. For the experiments in this thesis, pick-up forecasting with booking curve detruncation were used for this step. For the simple Q-forecasting example, the forecasted Q-bookings in each time frame, $fcst_{tf}$, are shown below in Figure 26, assuming the total number of forecasted Q-bookings to come is 125.

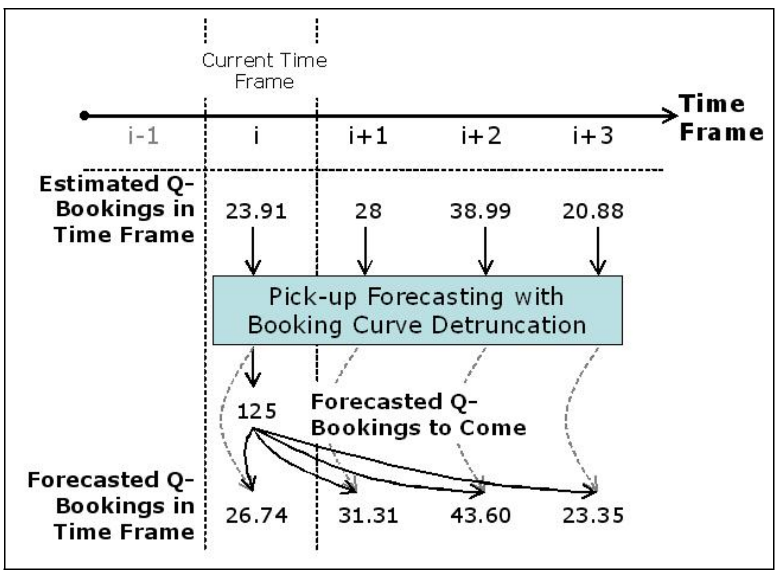

Figure 26: Example Forecasted Q-Bookings by Time Frame

The final step in forecasting potential demand is the partitioning of the total forecasted Q-bookings in each time frame into demand for the separate fare classes, using the formula below.

$$fcst_{f,tf} = fcst_{tf} \cdot \left(psup_{Q \to f-1,tf} - psup_{Q \to f,tf} \right)$$

Where:

- $fcst_{f,tf}$ signifies the mean forecasted demand in fare class f in time frame tf;
- $fcst_{tf}$ signifies the total forecasted equivalent demand for fare class Q in time frame tf;
- $psup_{Q \to f,tf}$ signifies the probability of sell-up from fare class Q to f (or f -1, the next higher fare class, as the case may be) in time frame tf;

For our example, the forecasted number of bookings in each time frame and fare class, $fcst_{f,,tf}$, is shown below in Figure 27. Note that no demand is forecasted for fare classes in time frames which must be closed due to AP requirements. Also, the total forecasted demand to come in each fare class is determined by summing over the remaining time frames.

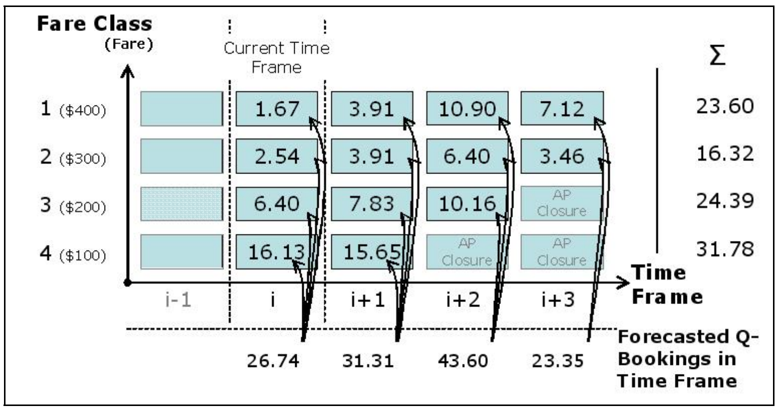

Figure 27: Example Forecasted Mean Demand by Time Frame and Fare Class

And at this point, the standard deviation of the forecasted demand in fare class f can also be calculated as shown in the formula below.

$$fcst\sigma_{f,tf} = fcst\sigma_{tf} \sqrt{psup_{Q \to f,tf} - psup_{Q \to f-1,tf}} = \frac{fcst\sigma_{tf}}{\sqrt{fcst_{tf}}} \sqrt{fcst_{f,tf}}$$

Where:

- $fcst\sigma_{tf}$ signifies the standard deviation of the forecasted Q-equivalent demand in time frame tf;

So at the end of this Q-forecasting process, we have a forecast of demand in each time frame and fare class based upon historical booking data. The process accounts for passenger sell-up in unrestricted fare structures, and produces mean and standard deviations values which allow for the use of the traditional seat optimizer methods.

Cléaz-Savoyen[71] and Vanhaverbeke[76] both used Q-forecasting within PODS, and described its potential for revenue improvements over traditional forecasting methods. But it is important to note that the method of Q-forecasting they used forecasted total bookings to come as well as sell-up to come, and did not account for AP restrictions (thus forecasting demand which can never be realized). The version used for this thesis represents a slight improvement. Besides differing by using time frame forecasting and sell-up, this version also accounts for AP restrictions – an enhancement which helps guard against over-forecasting of demand.

3.2.3 Hybrid Forecasting in PODS

At this point that we examine the combination of Section 3.1.3.2's traditional forecasting methodology (for fully-restricted fare structures) with Section 3.2.2's Q-forecasting (for unrestricted fare structures) to develop a "hybrid" forecast of demand for the PODS seat allocation optimizer in environments that are neither fully-restricted nor unrestricted. As discussed in Section 2.5.3, we assume that general passenger demand is actually a combination of product-oriented and price-oriented demand, and that it is important to differentiate between the two in order to avoid the spiral-down effect in semi-restricted fare structure environments.

Clearly, the biggest challenge in HF involves the identification these of product-oriented and price-oriented passengers, but using a separate forecasting methodology for each addresses this problem. Thus, the challenge becomes developing an accurate forecast for each demand segment based upon historical booking data.

It is not sufficient to simply classify all passengers who book in the lowest fare class available as price-oriented, because there are likely people who were specifically seeking that fare product and its restrictions. Conversely, we cannot blindly classify all passengers booking in the highest fare classes as product-oriented as the closure of lower fare classes may have enticed demand to sell-up from lower priced products.

But in developing ways to classify product-oriented and price-oriented demand, an airline is limited in its knowledge of the travel market. A given carrier clearly has full knowledge of its own fare structure and availability of each of its fare classes in a given OD market. But its knowledge of competitor offerings is often limited to awareness of a competitor's fare structure, but not seat availability beyond the lowest open class (and especially not historical booking data).

3.2.3.1 HF1, HF2, and HF3 Classifications

One way an airline can classify an historical booking as either product-oriented or price-oriented is by examining the other services available to the passenger at the time of booking. If multiple paths in an OD market are available, a carrier may be able to classify a passenger as product-oriented because that passenger chose the itinerary he or she did for a particular reason. In PODS, we have three different methods, or rules, for classifying bookings as product-oriented or price-oriented based exclusively on path availability; note that we define a path as closed (and unavailable) only when the highest fare class on that path is closed.

When a passenger books a ticket in a given fare class, he or she is counted as product-oriented if the next lower class is available…

- ...on the same path (i.e. same flight(s), same airline). This classification is known as "the path rule," or method **HF1**.
- …on some path provided by the same airline. This is known as the "the airline rule," or method **HF2**.
- …on some other path in the market (i.e. another airline). This is known as "the market rule," or method **HF3**.

All other bookings are classified as price-oriented, except as described in Section 3.2.3.2 below.

The general idea implicit in these three methods is that if a similar service is available in a cheaper fare class, any passenger that books the more expensive itinerary must be product-oriented. However, these three rules vary in defining similarity between paths from an origin to a destination.

3.2.3.2 Classifications IAP0, IAP1, IAP2

Beyond the three path availability methods, in PODS we also can classify historical bookings as product-oriented or price-oriented based upon the closure of lower classes – specifically closure due to advance purchase (AP) requirements (as opposed to closure by the seat allocation optimizer). We have three different methods for classifying bookings as product-oriented or price-oriented based exclusively on AP requirements of cheaper products.

In PODS, historical bookings where the next lower class is closed due to AP requirements are classified as…

- …price-oriented when using method **IAP0**.
- …product-oriented when using method **IAP1**.
- …product-oriented only if there is a difference in booking restrictions between the chosen class and the next lower fare class when using method **IAP2**.

3.2.3.3 Hybrid Forecasting Methodology

So PODS actually has nine versions of HF from which the user must choose (combining the three path availability methods and three advance purchase classifications). In the context of the PODS RM system for each airline, as shown in Figure 13, the hybrid forecaster first classifies historical bookings as either product-oriented or price-oriented. The price-oriented bookings are sent to the PODS Q-forecasting module, which forecasts bookings in each undifferentiated fare class. Likewise, the product-oriented bookings are sent to a traditional forecaster – pick-up forecasting is used in this thesis – which forecasts future product-oriented bookings in each fare class. The two sets of future bookings – product-oriented and price-oriented – are aggregated and sent to the PODS seat allocation optimizer. This methodology is shown below in Figure 28.

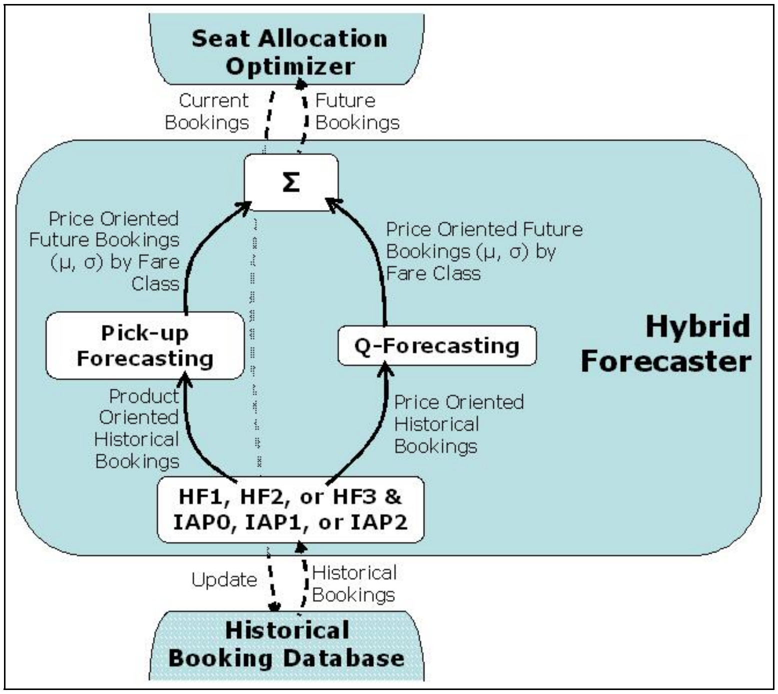

Figure 28: Hybrid Forecasting in PODS

3.2.4 Fare Adjustment

As mentioned earlier in Section 2.4.2, Fare Adjustment (FA) is a technique developed by Fiig and Isler[13] to improve airline revenues in unrestricted fare structure environments for carriers employing virtual nesting-based seat allocation optimizers. Because the simplified fare structures tested in this thesis contain undifferentiated fare classes, FA can be applied in the context of hybrid forecasting.

While Q-forecasting, as discussed above, deals with the class differentiation problem in the context of the RM system's demand forecaster, FA approaches the differentiation issue from within the seat allocation optimizer. Because an open class in an unrestricted fare structure must accept all demand, regardless of the WTP of those respective passengers, the opportunity to improve network revenues by inducing sell-up is unavailable.

In DAVN, FA logic lowers the pseudo fares used for optimizing booking limits in virtual nests (see Section 3.1.3.3) by an amount called the "Price Elasticity cost" or "PE cost," which reflects the lost opportunity for sell-up, as shown in the equation below. By reducing the pseudo fares in each undifferentiated fare structure by the PE cost, we can shift these fare classes into lower virtual nests, forcing these low revenue classes to close sooner and stimulating sell-up into the higher fare classes.

$$Pseudo\ Fare_f^{x,OD} = fare_f^{OD} - \sum_{\substack{j \in OD \\ x \neq j}} d_j - PE\ Cost^{OD}$$

This relationship can also be expressed as shown below.

$$Pseudo\ Fare_f^{x,OD} = fare'^{OD}_f - \sum_{\substack{j \in OD \\ x \neq j}} d_j$$

Where:

$$fare'^{OD}_f = fare^{OD}_f - PE\ Cost^{OD}$$

This value, *fare'*, also represents the expected marginal revenue the carrier can expect to receive once correcting for the lost sell-up revenue of the undifferentiated fare class, as shown in Figure 29.

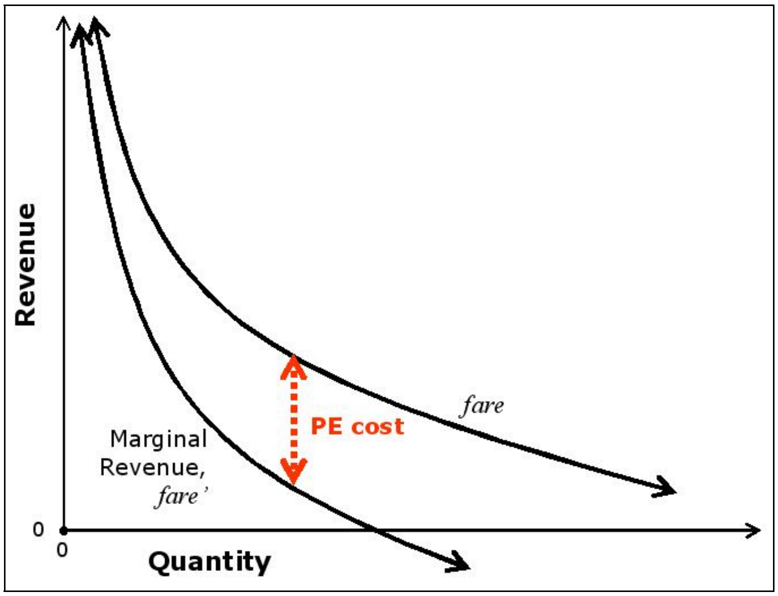

Figure 29: Relationship of Fare, Marginal Revenue (*fare'*), and PE cost

3.2.4.1 FA in PODS

Within PODS, we have two FA methods available: a continuous marginal revenue formulation (MR), and a discrete one (KI, for Karl Isler). The continuous FA method uses the formulation shown below to adjust fares.

$$fare'^{OD}_f = fare^{OD}_f - \frac{fare^{OD}_Q (FA\ FRAT5 - 1)}{-\ln(0.5)}$$

And the discrete fare adjustment formulation is shown below.

$$fare'^{OD}_f = \frac{psup_{Q \to f,tf} \cdot fare^{OD}_f - psup_{Q \to f-1,tf} \cdot fare^{OD}_{f-1}}{psup_{Q \to f,tf} - psup_{Q \to f-1,tf}}$$

In this thesis, we employ KI Fare Adjustment for all experiments involving FA. But for both of these methods, the sell-up notation is similar to that used in Section 3.2.2.2, with the exception of the FRAT5s.

3.2.4.2 Use of FA FRAT5 Values

Because FA seeks to capture sell-up behavior in unrestricted environments, some estimate of passenger WTP must be employed by the seat allocation optimizer. As passenger WTP increases, the PE cost used in FA must increase in order to close lower classes more quickly. Likewise, we would expect a lower PE cost to reflect lower WTP.

As discussed in Section 3.2.2.1, PODS makes use of FRAT5 values in Q-forecasting to model WTP. But as Cléaz-Savoyen[71] describes in this Master's thesis, the FRAT5 values used for fare adjustment must be lower than those used for Q-forecasting. To solve this problem, he uses a set of linear (i.e., not S-shaped) FA FRAT5s which are independent of the forecasting FRAT5 values.

But because each airline in PODS assumes WTP for its passengers, it is unrealistic to use two unrelated FRAT5 values for Q-forecasting and Fare Adjustment. In this thesis, we relate the two sets of FRAT5 values with a scaling factor, as shown below.

$$FA\ FRAT5_{tf} = 1 + f5scl\left(FRAT5_{tf} - 1\right)$$

Where:

- $FA\ FRAT5_{tf}$ signifies the FA FRAT5 value in a particular time frame;
- $FRAT5_{tf}$ signifies the FRAT5 value used for Q-forecasting in a particular time frame;
- $f5scl$ signifies the scaling factor, between 0 and 1, between the two sets of FRAT5s;

So for each airline within PODS employing Q-forecasting (or HF) and FA, the user must input not only a FRAT5 series which best describes passenger behavior for the forecaster, but also an appropriate scaling factor for Fare Adjustment. For example, an airline within PODS that assumes passenger sell-up behavior can best be modeled with FRAT5 series "C" (see Section 3.2.2.1) must also scale that set of forecasting FRAT5 values for use with FA. This example is illustrated in Figure 30 for scaling values between 0.1 and 0.5.

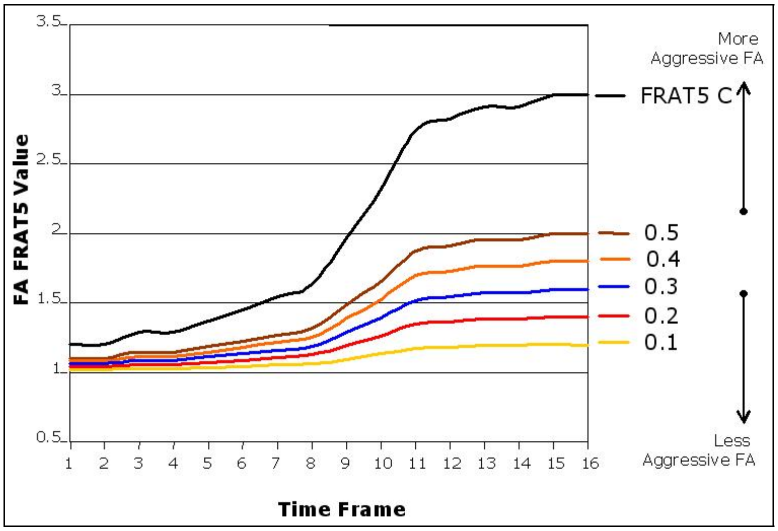

Figure 30: Selected PODS FA FRAT5 Values by Time Frame with Different Scaling Factors (FRAT5 "C")

3.2.5 Varying Sell-up by Path Quality in PODS

As briefly mentioned in Section 1.3.2, we intuitively expect that WTP is higher for non-stop service, when available, and that any airline offering superior (i.e. non-stop) service to the competition has the opportunity to extract more revenue in those specific OD markets. By examining each path available through an airline's particular network in the context of the overall air transportation network, we can selectively index all paths on all carriers in terms of relative path quality.

Within PODS, the capability exists to index each path into one of three distinct Path Categories defined by this relative path quality. Each path provided by each carrier can be labeled as…

- **…PCAT1** if the path is non-stop while the competition in the market only provides connecting service.
- **…PCAT2** if the path is non-stop and the competition provides non-stop service, or if the path is connecting and the competition provides only connecting service.
- **…PCAT3** if the path is connecting and the competition provides non-stop service.

In this thesis, we perform only experiments in which the behavior of PCAT1 is varied by assuming higher WTP (i.e. using higher FRAT5 values) – a practice referred to "Path Categorization" (PCAT). We test the expectation that an airline can enjoy revenue gains by exploiting the higher WTP for direct service through more aggressive HF and FA in its dominant OD markets.

3.3 Chapter Summary: RM in PODS

In this chapter, we have presented the Passenger Origin-Destination Simulator used to test Hybrid Forecasting in this thesis. Specifically, we described the two components of PODS: the Passenger Choice Model and the airline Revenue Management system. We then explained the evolution of RM techniques within PODS to address less-restricted fare structures, including Hybrid Forecasting.

In Chapter 4, we describe both the simulated air transportation network within PODS and the specific experiments performed in order to identify the impacts of categorizing product-oriented and price-oriented historical bookings for HF, given the limited booking information available to an airline; we also attempt to measure the value of Hybrid Forecasting in terms of the network revenue improvement for a particular airline. Chapter 5 describes attempts to improve the performance of HF by accounting for the effects of path quality on sell-up potential, as well as by using Fare Adjustment within an airline RM system.

4 Defining and Testing Hybrid Forecasting

In this chapter, we present and analyze a series of experiments using the Passenger Origin-Destination Simulator (PODS); our goal is to test Hybrid Forecasting (HF) as an effective Revenue Management (RM) tool in semi-restricted fare structures. For the simulations in this chapter (as well as in Chapter 5), we model an air transportation network served by two carriers: “Airline 1” (**AL1**) and “Airline 2” (**AL2**). This network is herein referred to as “Network D-6” and is described in Section 4.1. Note that all experiments are described in terms of “the RM system used by AL1 versus that used by AL2” in Network D-6.

In Section 4.2, we test all nine combinations of HF1, HF2, and HF3, and IAP0, IAP1, and IAP2, as described in Sections 3.2.3.1 and 3.2.3.2, to hypothesize which particular combination may be best suited for HF. We perform these tests with different estimates of passenger WTP (in terms of FRAT5 input series, see Section 3.2.2.1) for the airline utilizing HF (Airline 1) given that both carriers use the DAVN seat allocation optimizer; we also repeat these experiments in simulated environments of low and high demand, as well as for the EMSRb versus EMSRb case to test the use of HF with a leg-based optimizer. And we close this chapter by testing for the potential revenue gains of HF over pick-up forecasting, as well as examining the sensitivity of HF to various FRAT5 series in Section 4.3.

4.1 Experimentation Environment

In this section we introduce the simulated competitive environment used for all experiments in this thesis. Specifically, we present our simulated air transportation network and the carriers providing Origin to Destination (OD) service through it.

4.1.1 Description of Network D-6

Network D-6 is characterized by two competing hub-and-spoke carriers: AL1 and AL2. In the context of the US domestic air transportation, Airline 1’s hub is the centrally located Minneapolis-Saint Paul (MSP) International Airport and Airline 2’s hub is Dallas-Fort Worth (DFW) International Airport. The two carriers compete in a one-way, West to East network consisting of twenty Western spoke cities (Cities 1 through 20) and twenty Eastern spoke cities (Cities 21 through 40), not including the two hubs (H1 and H2, or Cities 41 and 42, respectively). Airline 1’s route network is shown in Figure 31, as is Airline 2’s in Figure 32.

Also, each airline operates three connecting banks of flights, in which flights arrive, simulated passengers change aircraft, and then the aircraft and passengers continue on to their final destinations. Both airlines have banks which begin (i.e., the aircraft arrive) at 10:30AM, 2:00PM, and 5:30PM.

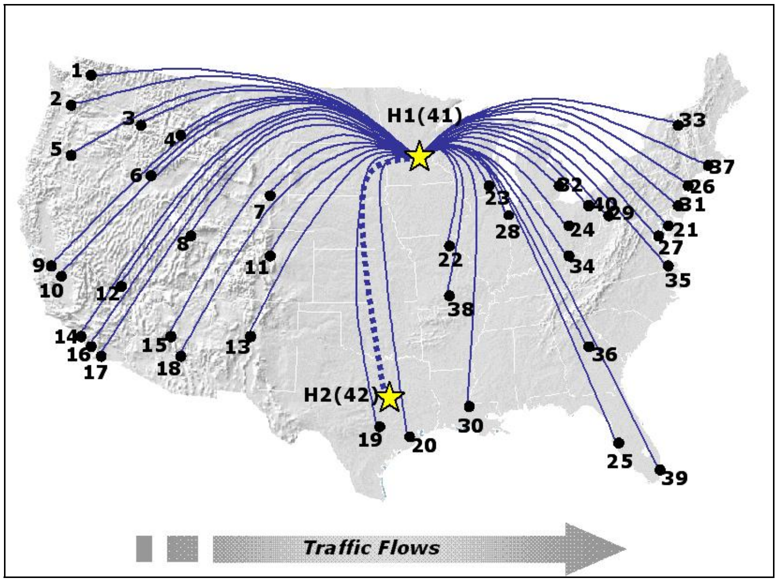

Figure 31: Route Network for Airline 1 in PODS Network D-6

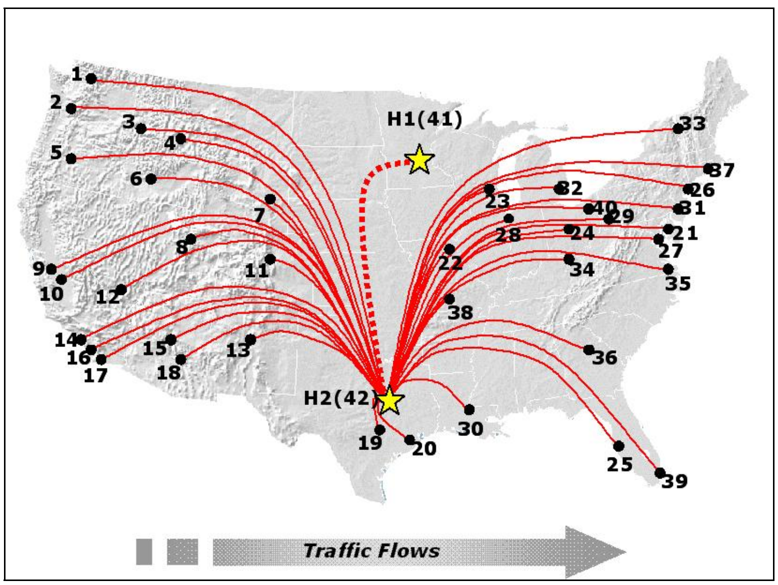

Figure 32: Route Network for Airline 2 in PODS Network D-6

As shown in these figures, each airline operates non-stop service between its hub and the 20 Western spokes, the 20 Eastern spokes, and the other airline's hub. Network D-6 is characterized by 252 total flight legs, as described below; note that the network consists of 2 airlines and 3 banks per airline.

- 20 legs per airline per bank from Cities 1 through 20 to the hub;
- 20 legs per airline per bank from the hub to Cities 21 through 40;
- 1 leg per airline per bank from the hub to the other airline's hub;
- 1 leg per airline per bank from the other airline's hub to the hub;

Network D-6 is also characterized by 2,892 total OD paths, as described below; again, note that the network consists of 2 airlines and 3 banks per airline.

- 400 paths per airline per bank from Cities 1 through 20 to Cities 21 through 40;
- 20 paths per airline per bank from Cities 1 through 20 to the hub;
- 20 paths per airline per bank from Cities 1 through 20 to the other airline's hub;
- 20 paths per airline per bank from the hub to Cities 21 through 40;
- 20 paths per airline per bank from the other airline's hub to Cities 21 through 40;
- 1 path per airline per bank from the hub to the other airline's hub;
- 1 path per airline per bank from the other airline's hub to the hub;

So for each airline and each bank, 482 OD paths exist. In terms of path quality, as described in Section 3.2.5, these paths can be classified as follows:

- 40 paths categorized as PCAT1, in which the airline provides superior path quality to the competitor;
- 402 paths categorized as PCAT2, in which both airlines provide equivalent service in terms of path quality;
- 40 paths categorized as PCAT3, in which the competitor airline offers superior path quality;

4.1.2 Semi-Restricted Fare Structures in Network D-6

As described in Section 2.3, increased competition from LCCs has led many legacy carriers to simplify their traditional fare structures by relaxing Saturday night stay restrictions, itinerary change fees, and non-refundability requirements, as well as by easing advance purchase (AP) requirements. The assignment of passenger disutilities for these restrictions is discussed in Section 3.1.2.2. Shown in Table 6 are examples of a traditional, fully-restricted fare structure as well as a parallel totally unrestricted fare structure for an airline offering six fare classes. Note that Fare Class 1 is the highest, most expensive class.

Table 6: Fully-restricted and Unrestricted Fare Structures in Network D-6

Fare Class	Fully-restricted AP	R1	R2	R3	Fare Class	Unrestricted AP	R1	R2	R3
1	0	0	0	0	1	0	0	0	0
2	3	0	1	0	2	0	0	0	0
3	7	1	0	0	3	0	0	0	0
4	10	1	1	0	4	0	0	0	0
5	14	1	1	1	5	0	0	0	0
6	21	1	1	1	6	0	0	0	0

In the above table, as well as within PODS, the following nomenclature is used:

- **AP** refers to advance purchase restrictions, or the number of days before departure in which the fare class will automatically be closed, independent of the seat allocation optimizer;

- **R1** refers to the traditional Saturday night stay restriction;
- **R2** refers to existence of a fee for a passenger changing his or her itinerary;
- **R3** refers to the non-refundability restriction;

A "1" indicates the presence of a certain restriction (R1, R2, or R3) in a fare structure, whereas a "0" indicates its absence. In the fully-restricted fare structure, the highest fare class (Fare Class 1), is traditionally intended for high yield business passengers and is characterized by an absence of any kind of booking restrictions. Indeed, nearly each fare class has a unique combination of booking restrictions traditionally used to fence passenger demand. But in the unrestricted fare structure, the six fare classes are identical with the exception of price, as discussed previously.

In this thesis, we examine RM in so-called "semi-restricted" fare structures, which lie somewhere between the fully-restricted and totally unrestricted structures of Table 6. The particular fare structure used by both airlines for all our experiments is shown in Table 7. As shown below, Fare Classes 3, 4, 5, and 6 share the same combination of booking restrictions (only R2 and R3), and AP requirements are used to proactively close these undifferentiated fare classes. In PODS, the fares charged by each airline for each fare product vary by OD path; the average fare charged in each fare class is shown in Table 8.

Table 7: Semi-Restricted Fare Structure

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

Table 8: Average Fares by Fare Class in PODS Network D-6

Fare Class	1	2	3	4	5	6
Average Fare	$412.85	$293.34	$179.01	$153.03	$127.05	$101.06

For consistency, Airline 1 is always the subject of experimentation, and its forecasting methodology is our variable of interest. For all experiments, Airline 1 uses the Displacement Adjusted Virtual Nesting (DAVN) seat optimizer in conjunction with various forms of HF (as well as Fare Adjustment and Path Categorization, in Chapter 5), unless otherwise noted. Its competitor, Airline 2, uses DAVN with conventional pick-up forecasting, unless otherwise noted. Because the goal of this thesis is to demonstrate improvements in airline RM due to HF, the primary statistic of interest is the percentage revenue increase (or loss) of Airline 1 due to a particular method.

4.1.3 Performance Measures and Base Cases in Network D-6

In an effort to explain changes in network revenue for each simulation, we examine the fare class mix of each airline, as well as other typical measures of general interest in the airline industry, including load factor and yield, as defined below.

- Load Factor (**LF**) is a measure of aircraft seat utilization, and is the ratio of passengers to seats for any given flight. Over a network, it is the ratio of the number of miles flown by a carrier's passengers, or revenue passenger miles (RPM), to the number of miles flown by the carrier's aircraft seats, or available seat miles (ASM).
- Yield is a measure of average revenue paid per passenger, normalized by distance flown. Over a network, it is the ratio of total revenue paid by passengers to the total RPMs.

In economic terms, the air transportation industry is characterized by extremely high fixed costs. For example, the cost for an airline to fly an aircraft on a single leg includes fuel, crew salaries, ownership expenses, etc. In contrast, the marginal cost of providing that air service to an additional passenger is very small (including baggage handling cost, on board services, reservation cost, etc.). For this reason, the marginal revenue of carrying an additional passenger (the fare paid) nearly always exceeds that passenger's marginal cost to the airline.

So unused seats, or idle capacity, are often avoided by airlines as they represent lost revenue opportunities. This is the reason for the traditional emphasis on high load factor within the airline industry, and its importance as a performance measure in a given carrier's network.

However, it is easy to envision a scenario where an airline compromises its overall health in an effort to maximize load factor. As a given carrier continues to lower fares in order to attract more passengers, the average fare paid by per passenger will decrease – the case of dilution, as originally shown in Figure 2 in Chapter 1.

For this reason, passenger yield is another critical metric for gauging the performance of an airline's RM system. But similar to the exercise of maximizing load factor, a carrier seeking to maximize yield by increasing its fares will quickly find that it denies too many bookings and again compromises its total revenue. This is the case of overprotection, as also shown in Chapter 1's Figure 2.

Therefore, an airline seeking to maximize its revenue (not load factor or yield) will exhibit some "healthy" combination of the two, and it is important to study all three metrics – revenue, LF, and yield – together in order to appraise a carrier's RM system.

To properly evaluate HF's potential revenue contributions to a given airline, we present several "Base Cases" (**BC**) which represent appropriate control environments. These Base Cases act as the "before" scenarios for comparison to the "after" effects of our HF experiments.

For the first Base Case (BC1), both Airline 1 (AL1) and Airline 2 (AL2) use identical Revenue Management systems: DAVN optimization with standard pick-up forecasting. The results of this simulation are shown in Table 9.

Table 9: Base Case 1 Results – DAVN versus DAVN

Airline	Revenue	Load Factor (%)	Yield ($/RPM)
AL1	$1,040,277	82.46	0.1028
AL2	$1,030,550	82.25	0.0984

In this Base Case, Airline 1 enjoys approximately 1% more network revenue than its competitor due to asymmetry within Network D-6. Its load factor and yield are also higher than Airline 2. For the experiments in this thesis, we are concerned with the possible revenue improvements for Airline 1 when using various forms of HF.

Furthermore, because the efficacy of Hybrid Forecasting's price-oriented component – Q-forecasting – is governed by passenger sell-up from lower to higher fare classes, the estimation of passenger WTP is critical. As discussed in Section 3.2.2.1, we use "FRAT5" values in PODS as a proxy of WTP, and we have 17 pre-defined FRAT5 series from which a simulated airline can choose to represent WTP of its passenger.

To simplify analysis, as well as to ease the simulation effort of PODS, most experiments in this chapter and Chapter 5 use FRAT5 series "C" as the default "baseline" estimate of WTP. This is because we believe series "C" represents a reasonable estimate of WTP, and is moderate within the range of WTP estimates available (i.e. not too aggressive or too passive with respect to sell-up).

4.2 Defining Hybrid Forecasting: HF1, HF2, HF3 and IAP0, IAP1, IAP2

The first step in improving a given airline's revenue with Hybrid Forecasting is to define what exactly constitutes a product-oriented historical booking and a price-oriented one. After doing so, traditional pick-up forecasting and Q-forecasting can be used to estimate future product-oriented and price-oriented bookings, respectively.

As described in Section 3.2.3.1, we have three ways to classify historical bookings in terms of path availability: the path rule (HF1), the airline rule (HF2), and the market rule (HF3). And we also have three ways to classify historical bookings where the next lowest class was closed due to Advance Purchase (AP) requirements: IAP0, IAP1, and IAP2, as described in Section 3.2.3.2 above. In the following sections, we repeatedly experiment with all nine of these combinations to determine the "best" performing of the nine definitions of product-oriented and price-oriented demand.

4.2.1 <u>DAVN w/ HF (FRAT5 "C") versus DAVN</u>

In this first series of simulations, Airline 1 uses DAVN with HF against a competitor using DAVN with pick-up forecasting. All nine possible definitions of product-oriented and price-oriented demand were tested, and the revenue results for AL1 are shown in Figure 33. Because the price-oriented component of HF – Q-forecasting – requires an estimate of passenger Willingness-to-Pay (WTP), as described in Section 3.2.2 above in order to assess sell-up, we have selected our baseline FRAT5 Series of "C", as described in Section 4.1.3.

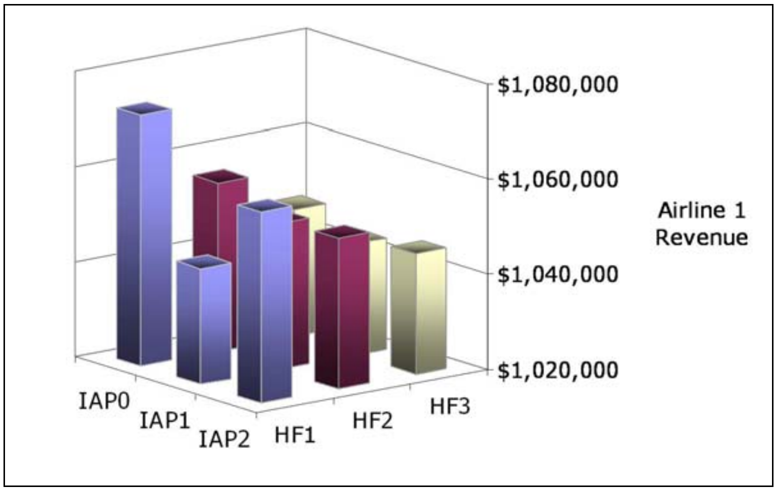

Figure 33: Airline 1 Revenues, DAVN w/ HF versus DAVN

As shown in Figure 34 and Table 10, the combination of HF1 and IAP0 leads to the largest revenue increase, by far, for Airline 1 over BC1 – 3.11%, or $1,072,646 to $1,040,277 in the Base Case.

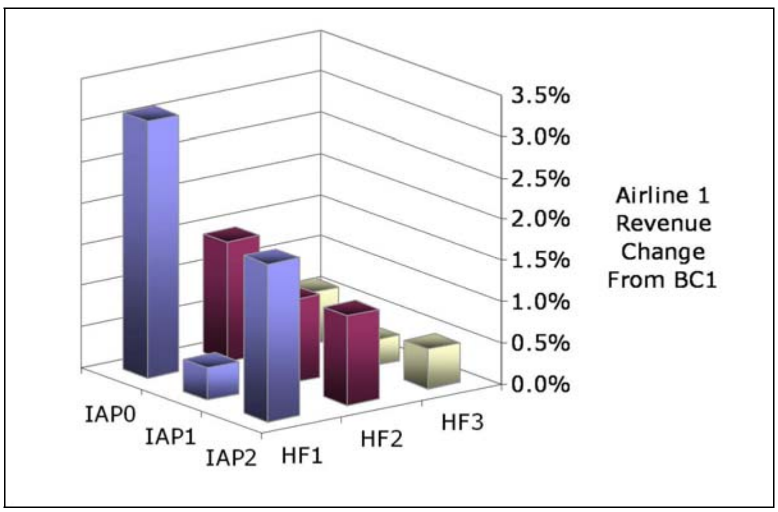

Figure 34: Airline 1 Revenue Changes from BC1, DAVN w/ HF versus DAVN

Table 10: DAVN versus DAVN, Airline 1 Revenue and Change from BC1

DAVN vs. DAVN	IAP0		IAP1		IAP2	
	Revenue	Δ	Revenue	Δ	Revenue	Δ
HF1	$1,072,646	3.11%	$1,044,303	0.39%	$1,060,161	1.91%
HF2	$1,055,332	1.45%	$1,050,673	1.00%	$1,051,538	1.08%
HF3	$1,046,757	0.62%	$1,043,464	0.31%	$1,045,284	0.48%

To focus on the effects of the individual path availability rules and AP purchase classifications, we center on this HF1-IAP0 combination and incrementally refine the product-oriented/price-oriented definition in the following sections.

4.2.1.1 HF1, HF2, HF3

By comparing HF1-IAP0, HF2-IAP0, and HF3-IAP0, we can isolate the effects of the HF1, HF2, and HF3 booking classification definitions. The cumulative bookings in each of the six fare classes, or the fare class mix, are shown below in Figure 35, Figure 36, and Figure 37, for HF1, HF2, and HF3, respectively.

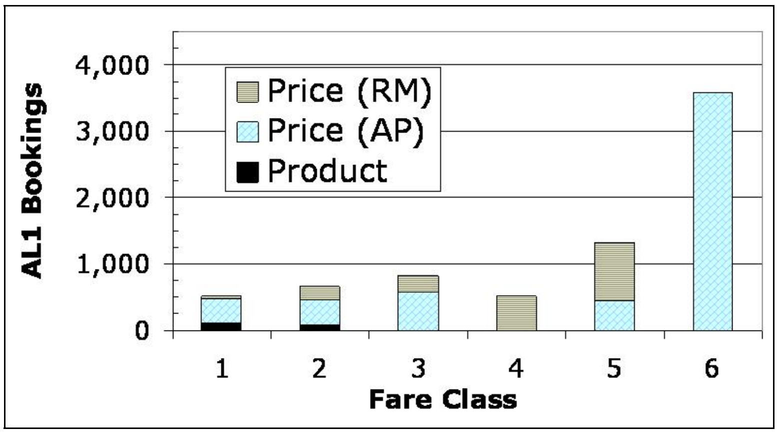

Figure 35: Fare Class Mix for HF1-IAP0

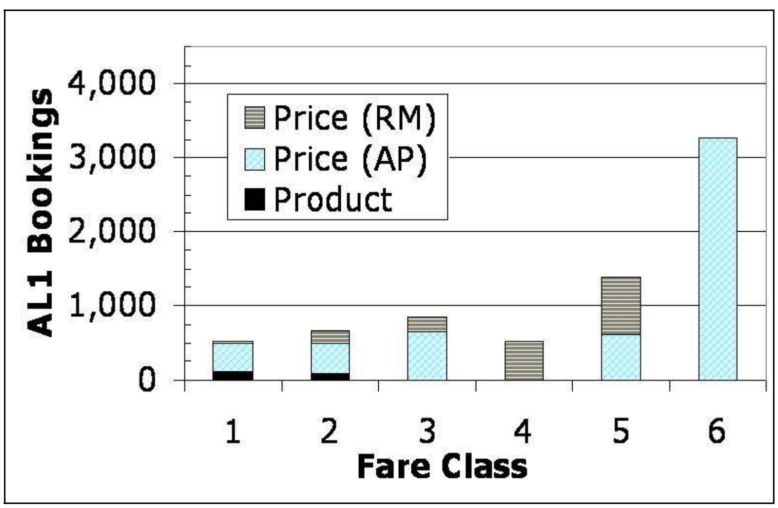

Figure 36: Fare Class Mix for HF2-IAP0

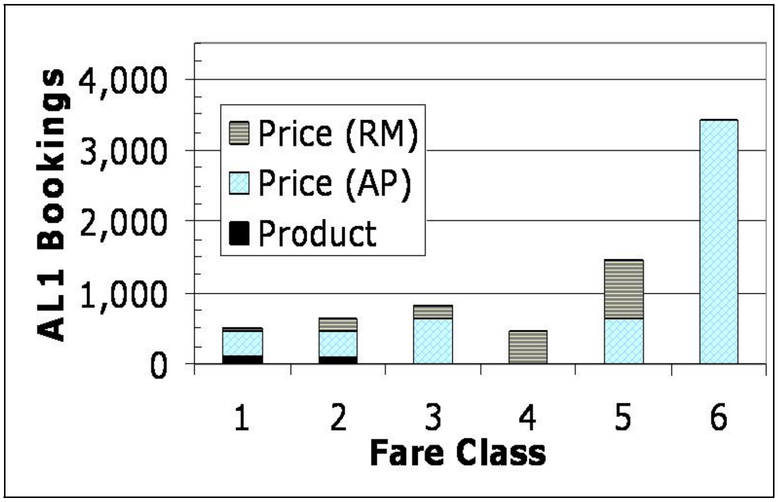

Figure 37: Fare Class Mix for HF3-IAP0

Illustrated in these Figures, as well as in Table 11, are the proportions of product-oriented and price-oriented bookings in each fare class; furthermore, the price-oriented bookings are distinguished between those where the next lower class is closed by the RM system's seat allocation optimizer (RM) and those automatically closed by advance purchase (AP) requirements.

Table 11: Product-oriented and Price-oriented Bookings by Fare Class – HF1, HF2, HF3

Fare Class	HF1			HF2			HF3		
	Product	Price		Product	Price		Product	Price	
		AP	RM		AP	RM		AP	RM
1	117	357	41	120	376	28	116	356	30
2	79	388	194	85	401	170	83	387	160
3	0	584	233	0	653	195	0	648	167
4	0	0	516	0	0	530	0	0	469
5	0	441	884	0	608	783	0	644	810
6	0	3,591	0	0	3,257	0	0	3,409	0
Σ	197	5,362	1,867	205	5,295	1,705	199	5,444	1,636

As shown in Table 11, these three scenarios actually have very similar fare class mixes. The numbers of product-oriented bookings are 197, 205, and 199 – a tight spread. Of significance is that HF1-IAP0 has the most bookings in the lower classes. In Fare Classes 4, 5, and 6, this combination has 5,431 total bookings, compared to 5,178 for HF2-IAP0 and 5,332 for HF3-IAP0. We hypothesize that these extra bookings in the lower fare classes somehow drive HF1's revenue above the HF2 and HF3 booking classification definitions – a hypothesis we further investigate by examining the operating measures for each case.

The revenue, load factor, and yield for both airlines in each of these three scenarios are shown in Table 12. The HF1-IAP0 combination for Airline 1 not only has the highest revenue of the three (as well as of the nine total combinations), but also has the highest LF, by far, and the lowest yield. The coupling of high revenues and loads for HF1 implies that HF2 and HF3 over-forecast demand in the higher fare classes.

Doing so likely leads the seat optimizer to overprotect in these classes (see Figure 2) and lose bookings in the lower fare classes.

Table 12: Results for HF1, HF2, HF3 & IAP0, DAVN w/ HF versus DAVN

Operating Statistic	Airline	HF1-IAP0	HF2-IAP0	HF3-IAP0
Revenue	AL1	$1,072,646	$1,055,332	$1,046,757
	AL2	$1,014,617	$1,030,502	$1,029,679
Load Factor (%)	AL1	86.20	82.04	82.86
	AL2	81.42	82.85	82.47
Yield ($/RPM)	AL1	0.1014	0.1049	0.1030
	AL2	0.0978	0.0976	0.0980

Thus, it appears that HF1, the path rule, outperforms the other two definitions of product-oriented demand in terms of path availability at the time of booking. In our experiments here, the path rule does not seem to overprotect the higher fare classes as much as the market and airlines rules. Furthermore, the higher number of bookings in fare classes 4, 5, and 6 (as well as overall LF) for HF1-IAP0 supports the idea that HF1 outperforms HF2 and HF3.

4.2.1.2 IAP0, IAP1, IAP2

By comparing HF1-IAP0, HF1-IAP1, and HF1-IAP2, we can isolate the effects of the IAP0, IAP1, and IAP2 AP classification definitions. The fare class mixes are shown below in Figure 38, Figure 39, and Figure 40, for IAP0, IAP1, and IAP2, respectively.

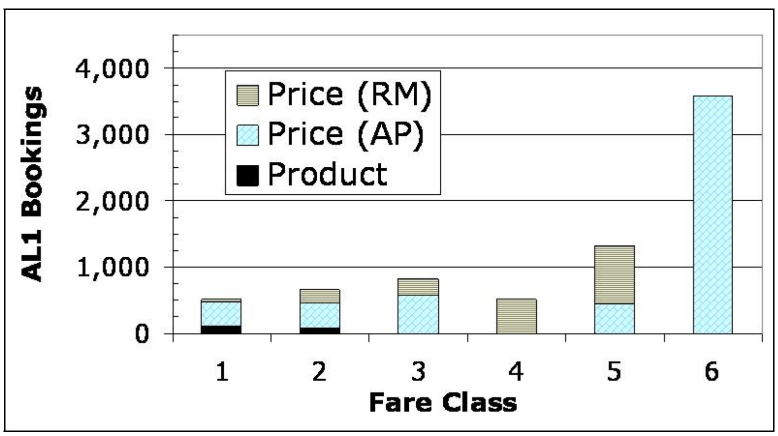

Figure 38: Fare Class Mix for HF1-IAP0

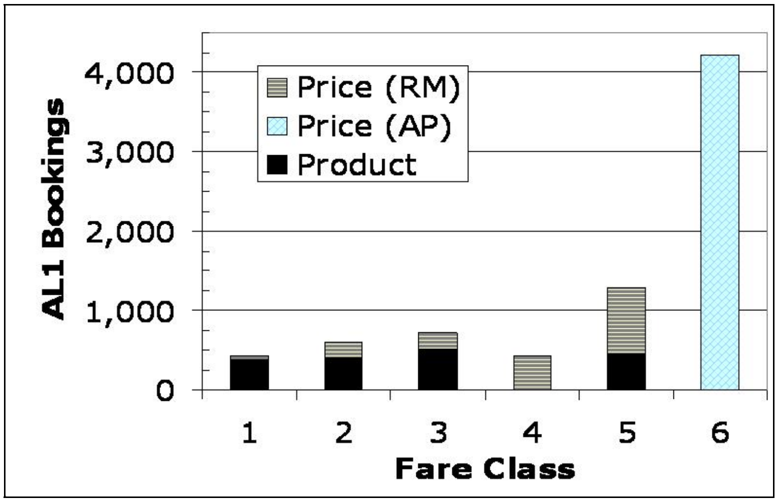

Figure 39: Fare Class Mix for HF1-IAP1

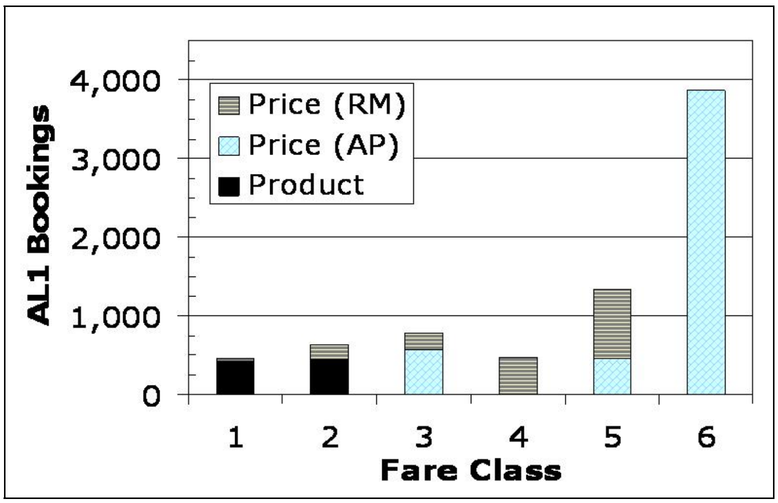

Figure 40: Fare Class Mix for HF1-IAP2

As with our simulations of the three path availability rules, the proportions of product-oriented and price-oriented bookings (both AP and RM) by fare class with IAP0, IAP1, and IAP2 are presented in these Figures, as well as in Table 13. But relative to our comparison across HF1, HF2, and HF3, there is much more variation in booking levels across the three AP methods.

Table 13: Product-oriented and Price-oriented Bookings by Fare Class - IAP0, IAP1, IAP2

Fare Class	IAP0			IAP1			IAP2		
	Product	Price		Product	Price		Product	Price	
		AP	RM		AP	RM		AP	RM
1	117	357	41	381	0	51	423	0	42
2	79	388	194	401	0	189	446	0	186
3	0	584	233	518	0	197	0	570	217
4	0	0	516	0	0	435	0	0	468
5	0	441	884	449	0	833	0	457	875
6	0	3,591	0	0	4,210	0	0	3,867	0
Σ	197	5,362	1,867	1,749	4,210	1,705	869	4,894	1,789

Upon examination of the fare class mix for each of these simulations, we can make the following observations about IAP0, IAP1, and IAP2:

- Based upon the definitions in Section 3.2.3.2, it is impossible to observe any AP price-oriented bookings in Fare Class 4 under any of the three classifications because Fare Classes 4 and 5 close simultaneously 14 days before departure (see semi-restricted fare structure in Section 4.1.2).
- IAP0 shows the most AP price-oriented bookings and the fewest product-oriented bookings. This is a result of its blind labeling of all historical bookings in which the next lowest class closed due to AP requirements as price-oriented.
- Conversely, IAP1 shows most product-oriented bookings and the fewest AP price-oriented bookings because it blindly labels so many product-oriented bookings, and IAP2 falls between IAP0 and IAP1.

The revenue, load factor, and yield for both airlines in each of these three scenarios are shown in Table 14. The HF1-IAP0 combination for Airline 1 not only has the highest revenue of the three (as well as of the nine total combinations), but also has the lowest LF, by far, and the highest yield (in direct contrast with our simulations across HF1, HF2, and HF3). The coupling of higher revenues and high yields for IAP0 implies that IAP1 and IAP2 under-forecast demand in the higher fare classes. Doing so likely leads the seat optimizer to dilute revenue with excess bookings in the lower fare classes (see Figure 2) and lose the revenue of high yield passengers.

Table 14: Results for HF1 & IAP0, IAP1, IAP2, DAVN w/ HF versus DAVN

Operating Statistic	Airline	HF1-IAP0	HF1-IAP1	HF1-IAP2
Revenue	AL1	$1,072,646	$1,044,303	$1,060,161
	AL2	$1,014,617	$1,030,362	$1,017,538
Load Factor (%)	AL1	86.20	89.10	88.01
	AL2	81.42	80.10	80.75
Yield ($/RPM)	AL1	0.1014	0.0955	0.0982
	AL2	0.0978	0.1010	0.0989

This hypothesis is supported by the larger number of product-oriented bookings with IAP1 and IAP2 – bookings which increase the use of traditional pick-up forecasting.

Because it is this forecasting method (when used inappropriately in less restricted fare structures) which accelerates the revenue dilution of the spiral-down effect (see Section 2.3.2), the heavier use of Q-forecasting for the price-oriented demand in HF1-IAP0 implies IAP0 is characterized by less spiral-down than IAP1 and IAP2.

4.2.1.3 Summary of DAVN w/ HF (FRAT5 "C") versus DAVN

Based upon these simulations of HF using FRAT5 series "C", the HF1-IAP0 combination outperforms the other eight in terms of network revenue improvement for Airline 1 over Base Case 1. The higher revenue and LF observed with the path rule (HF1) compared to HF2 and HF3 indicate that the latter two rules overprotect the higher fare classes and lose out on revenue gains from selling in lower fare classes. And the higher revenue and yields observed with IAP0 compared to IAP1 and IAP2 indicate that the latter two classification methods are characterized by more spiral-down and revenue dilution. Thus, HF1-IAP0 represents the "best" combination of load factor and yield, and leads to the best revenue improvement over standard pick-up forecasting.

4.2.2 Low and High Demand

To test the robustness of the HF1-IAP0 combination, we first examine the performance of all nine HF definitions under cases of lower and higher demand. More specifically, the experiments performed in Section 4.2.1 used a Demand Multiplier (DM) of 0.9, as previously mentioned in Section 3.1.2.1. Here we use a DM of 0.8 to simulate lower demand and a DM of 1.0 to simulate higher demand.

4.2.2.1 New Base Cases

Because Base Case 1 uses a baseline DM of 0.9, we must simulate two more BCs for low and high demand. Base Case 2 is identical to Base Case 1, except with an overall DM of 0.8. The results of this simulation are shown in Table 17. Due to the lower demand, both AL1 and AL2 experience lower network revenues, load factors, and yields compared to BC1.

Table 15: Base Case 2 Results – DAVN versus DAVN, Low Demand

Airline	Revenue	Load Factor (%)	Yield ($/RPM)
AL1	$931,999	78.07	0.0973
AL2	$920,372	77.28	0.0935

In contrast, both airlines see higher network revenues, load factors, and yields in Base Case 3, as shown in Table 16; this simulation is identical to BC1 and BC2, expect the DM is 1.0.

Table 16: Base Case 3 Results – DAVN versus DAVN, High Demand

Airline	Revenue	Load Factor (%)	Yield ($/RPM)
AL1	$1,149,712	85.41	0.1097
AL2	$1,138,380	85.18	0.1049

As with BC1, Airline 1's network revenue exceeds that of its competitor in low and high demand environments. Also, Airline 1's load factors and yields exceed that of Airline 2 in both BC2 and BC3. These differences are likely due to asymmetries in Network D-6.

4.2.2.2 DAVN w/ HF versus DAVN, Low Demand

The changes in Airline 1's revenue when using the nine various HF methodologies in low demand environments are shown in Figure 41. As with the case of DM = 0.9, the HF1-IAP0 combination outperforms all others in terms of network revenue.

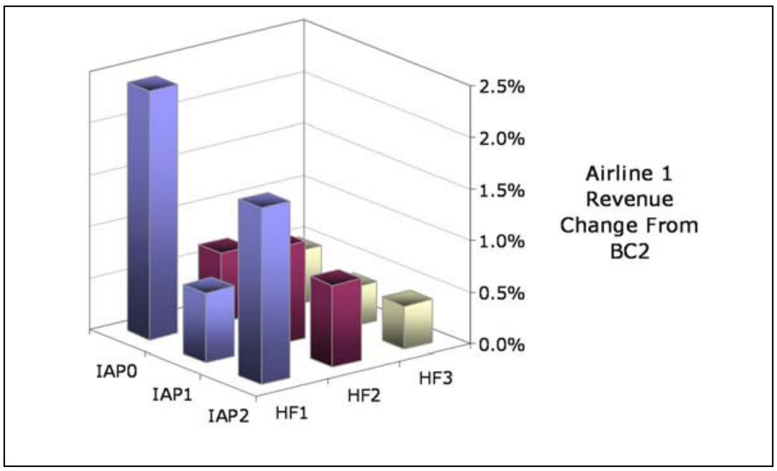

Figure 41: Airline 1 Revenue Changes from BC2, DAVN w/ HF versus DAVN for Low Demand

As shown in Table 17, HF1-IAP0 produces a revenue improvement of 2.41% over Base Case 2 – by far the largest increase of the nine historical booking classifications. And as with the baseline demand case, all nine definitions improve revenue to some extent.

Table 17: DAVN w/ HF versus DAVN for Low Demand, Airline 1 Revenue and Change from BC2

DM = 0.8	IAP0		IAP1		IAP2	
	Revenue	Δ	Revenue	Δ	Revenue	Δ
HF1	$954,416	2.41%	$938,194	0.66%	$947,869	1.70%
HF2	$938,253	0.67%	$940,705	0.93%	$939,330	0.79%
HF3	$936,753	0.51%	$935,467	0.37%	$935,911	0.42%

HF1 produces higher load factors than either HF2 or HF3 across a fixed IAP classification, as shown in Table 18. Furthermore, IAP0 always has higher yields than IAP1 and IAP2. These observations are consistent with the results for the baseline demand case.

Table 18: Airline 1 Load Factors and Yields for Low Demand

DM = 0.8	IAP0		IAP1		IAP2	
	LF (%)	Yield ($/RPM)	LF (%)	Yield ($/RPM)	LF (%)	Yield ($/RPM)
HF1	81.84	0.0951	83.73	0.0913	82.91	0.0932
HF2	77.39	0.0988	80.73	0.0950	78.74	0.0972
HF3	78.91	0.0968	79.40	0.0960	79.15	0.0964

So it appears that in low demand situations, HF2 and HF3 again lead to overprotection compared to HF1, and that IAP1 and IAP2 lead to more spiral-down of revenues than IAP0 – just as with the baseline DM of 0.9. So as with the baseline demand case, the HF1-IAP0 combination outperforms the other eight definitions of product-oriented and price-oriented historical bookings for use with Hybrid Forecasting. Of note is that the revenue gained over the appropriate Base Case for low demand – 2.41% - is smaller than the revenue gained for normal demand – 3.11%. This observation suggests that the benefits of HF are not independent of the level of passenger demand – an expected result because environments of low demand typically diminish the impact of RM systems.

4.2.2.3 DAVN w/ HF versus DAVN, High Demand

The changes in Airline 1's revenue when using the nine various historical booking definitions in high demand environments are shown in Figure 42. As with the cases of baseline and low demand (or DM = 0.9 and 0.8, respectively), the HF1-IAP0 combination outperforms all others.

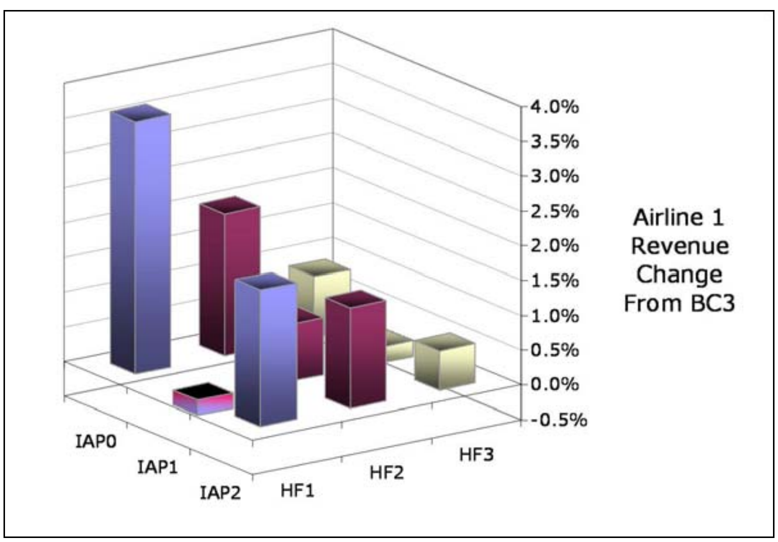

Figure 42: Airline 1 Revenue Changes from BC3, DAVN w/ HF versus DAVN for High Demand

As shown in Table 19, HF1-IAP0 produces a revenue improvement of 3.63% over Base Case 2 – again the largest increase. But unlike the normal and low demand cases, the HF1-IAP1 combination actually decreases Airline 1's revenue by 0.22% compared to Base Case 3. This fact suggests that the detrimental revenue effects of spiral-down can be intense enough in periods of high demand to actually hurt the airline more than if standard pick-up forecasting had been used instead – a critical observation which both demonstrates the potential severity of spiral-down, as well as provides strong evidence suggesting IAP1-HF1 combination is seriously flawed.

Table 19: DAVN w/ HF versus DAVN for High Demand, Airline 1 Revenue and Change from BC3

DM = 1.0	IAP0		IAP1		IAP2	
	Revenue	Δ	Revenue	Δ	Revenue	Δ
HF1	$1,191,395	3.63%	$1,147,195	-0.22%	$1,172,254	1.96%
HF2	$1,173,228	2.05%	$1,159,153	0.82%	$1,166,297	1.44%
HF3	$1,159,782	0.88%	$1,152,300	0.23%	$1,156,197	0.56%

HF1 produces higher load factors than either HF2 or HF3 across a fixed IAP classification, as shown in Table 20. Furthermore, IAP0 always has higher yields than IAP1 and IAP2. These observations are consistent with the results for baseline demand and low demand. They also suggest, once again, that the HF1 definition leads to less overprotection in high demand environments than HF2 and HF3, and IAP0 appears to guard against spiral-down better than IAP1 and IAP2.

Table 20: Airline 1 Load Factor and Yield for High Demand

DM = 1.0	IAP0		IAP1		IAP2	
	LF (%)	Yield ($/RPM)	LF (%)	Yield ($/RPM)	LF (%)	Yield ($/RPM)
HF1	88.74	0.1094	92.47	0.1011	91.16	0.1048
HF2	84.80	0.1128	88.78	0.1064	86.88	0.1094
HF3	85.44	0.1106	86.51	0.1086	86.04	0.1095

And just as the revenue improvement grew as demand increased from low to baseline levels, the revenue gained over the appropriate Base Case for high demand – 3.63% - is larger than the revenue gained for baseline demand – 3.11%. But as mentioned in Section 4.2.2.2 for low demand, this result is not surprising. RM systems provide disproportionate gains in high demand when booking requests are intensified. So not only does it appear that HF1-IAP0 again constitutes the best Hybrid Forecasting revenue improvement for an Airline using DAVN, we have observed that the percentage revenue gained over standard pick-up forecasting increases with the intensity of demand.

4.2.3 EMSRb w/ HF versus EMSRb

To further test the nine historical booking classifications, we simulate Hybrid Forecasting with a leg based seat allocation optimizer – the EMSRb heuristic, as described in Section 3.1.3.3. This is in contrast to the Origin-Destination (OD) path-based optimizer of DAVN used for all other experiments. For these simulations (and all others outside of Section 4.2.2) we again use the baseline Demand Multiplier of 0.9.

For the fourth Base Case (BC4), both airlines use identical Revenue Management (RM) systems: EMSRb optimization with standard pick-up forecasting. The results of this simulation are shown in Table 21. Due to the less sophisticated RM system, both AL1 and AL2 experience lower network revenues, load factors, and yields compared to BC1.

Table 21: Base Case 4 Results - EMSRb versus EMSRb

Airline	Revenue	Load Factor (%)	Yield ($/RPM)
AL1	$1,029,823	83.61	0.1004
AL2	$1,025,096	83.2	0.0967

In this Base Case, Airline 1 sees approximately 0.5% more network revenue than its competitor (a difference smaller than in BC1). Airline 1's load factor and yield is also less than in Base Case 1, though still superior to Airline 2's. These observations suggest that the network asymmetries inherent in Network D-6 exacerbate revenue gains for Airline 1 in the more sophisticated DAVN optimizer more than with EMSRb.

The changes in Airline 1's revenue when using the nine various HF methodologies in with an EMSRb seat optimizer are shown in Figure 43. As with the previous three cases (DAVN versus DAVN for baseline, low, and high demand), the HF1-IAP0 combination outperforms all others.

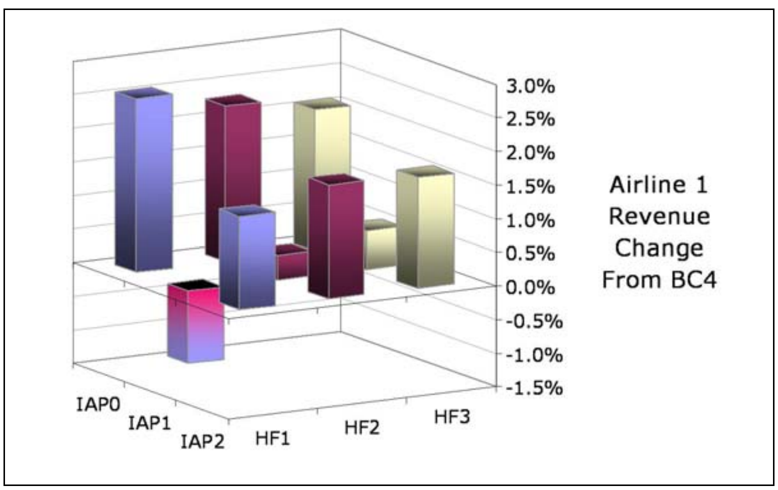

Figure 43: Airline 1 Revenue Changes from BC4, EMSRb w/ HF versus EMSRb

As shown in Table 22, HF1-IAP0 produces a revenue improvement of 2.58% over Base Case 4 – a slightly larger increase than that with HF2-IAP0 or HF3-IAP0. And as in the case of DAVN with HF versus DAVN, the HF1-IAP1 combination decreases Airline 1's revenue by 1.03% compared to Base Case 4 – a significant degree of spiral-down.

Table 22: EMSRb w/ HF versus EMSRb, Airline 1 Revenue and Change from BC4

EMSRb	IAP0		IAP1		IAP2	
	Revenue	Δ	Revenue	Δ	Revenue	Δ
HF1	$1,056,385	2.58%	$1,018,703	-1.08%	$1,044,050	1.38%
HF2	$1,053,652	2.31%	$1,033,670	0.37%	$1,047,265	1.69%
HF3	$1,051,456	2.10%	$1,035,775	0.58%	$1,046,743	1.64%

Regarding variations across HF1, HF2, and HF3, or IAP0, IAP1, and IAP2, the results are similar to those with the DAVN optimizer. HF1 produces higher load factors than either HF2 or HF3 across a fixed IAP classification, as shown in Table 23 – a clue that it leads to less overprotection due to its higher revenues. Furthermore, IAP0 always has higher yields than IAP1 and IAP2 – a clue that it guards against spiral-down better than the other two booking definitions.

Table 23: Airline 1 Load Factor and Yield for EMSRb w/ HF

EMSRb	IAP0		IAP1		IAP2	
	LF (%)	Yield ($/RPM)	LF (%)	Yield ($/RPM)	LF (%)	Yield ($/RPM)
HF1	84.35	0.1021	87.57	0.0948	86.16	0.0988
HF2	81.49	0.1054	84.58	0.0996	83.05	0.1028
HF3	80.25	0.1068	83.27	0.1014	81.76	0.1044

It is noteworthy that Airline 1's revenue improvement for HF1-IAP0 Hybrid Forecasting when using EMSRb – 2.58% - is smaller than that when using DAVN – 3.11%. So not only do we again observe that HF1-IAP0 outperforms the other 8 historical booking definitions for HF, but it also appears that the relative revenue gain over standard pick-up forecasting increases with the sophistication of the seat inventory optimization method. This latter observation is not a surprise, as the path-based DAVN optimizer should outperform the leg-based EMSRb in Network D-6.

4.2.4 DAVN w/ HF (FRAT5 "A4") versus DAVN

As previously mentioned in Section 4.1.3, we have been using the baseline FRAT5 Series "C" as Airline 1's moderate estimate of passenger WTP. In order to test the sensitivity of the nine historical booking definitions to the WTP estimate, we re-simulate DAVN with HF versus DAVN with a different FRAT5 series.

FRAT5 "A4" is more aggressive than "C", as previously described in Section 3.2.2.1. Thus, if a particular airline uses "A4", it assumes it can capture more sell-up by getting its passengers into higher fare classes than if that same airline uses "C". The changes in Airline 1's network revenue when using the nine various HF booking classifications and FRAT5 "A4" are shown in Figure 44; the appropriate Base Case for comparison is BC1. As with the case of FRAT5 "C", the HF1-IAP0 combination outperforms all others in terms of network revenue over BC1.

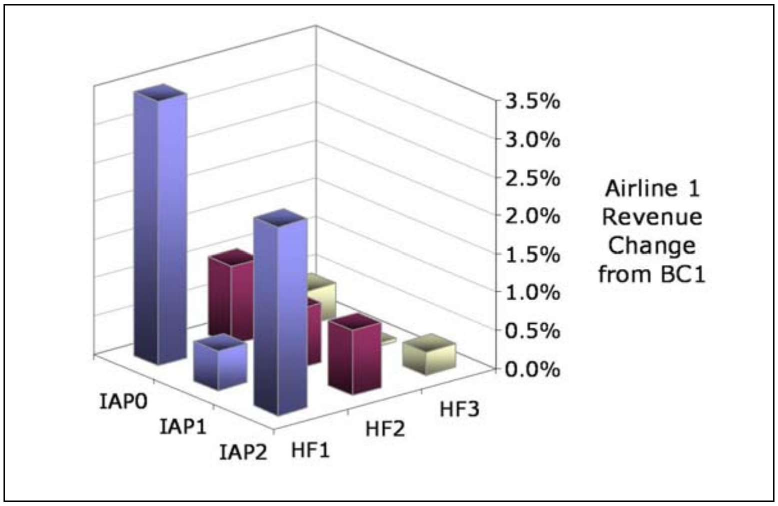

Figure 44: Airline 1 Revenue Changes from BC1, DAVN w/ HF versus DAVN (FRAT5 "A4")

As shown in Table 24, HF1-IAP0 produces a revenue improvement of 3.45% over Base Case 1 – by far the largest increase of the nine methods. And all methodologies improve network revenue for Airline 1 to some extent (though the HF3-IAP1 combination only improves AL1's situation by a paltry $483.

Table 24: DAVN w/ HF versus DAVN (FRAT5 "A4"), Airline 1 Revenue and Change from BC1

FRAT5 "A4"	IAP0		IAP1		IAP2	
	Revenue	Δ	Revenue	Δ	Revenue	Δ
HF1	$1,076,200	3.45%	$1,045,611	0.51%	$1,065,992	2.47%
HF2	$1,051,019	1.03%	$1,048,364	0.78%	$1,049,161	0.85%
HF3	$1,044,826	0.44%	$1,040,760	0.05%	$1,043,434	0.30%

As with every other scenario discussed in Section 4.2 (baseline demand, low demand, etc.) variations across HF1, HF2, and HF3, or IAP0, IAP1, and IAP2 with FRAT5 "A4" suggest reasons why HF1-IAP0 outperforms the other eight definitions. Once again, HF1 leads to higher load factors than either HF2 or HF3 across a fixed IAP classification, as shown in Table 25, and IAP0 always has higher yields than IAP1 and IAP2.

Table 25: Airline 1 Load Factor and Yield for DAVN w/ HF (FRAT5 "A4")

FRAT5 "A4"	IAP0		IAP1		IAP2	
	LF (%)	Yield ($/RPM)	LF (%)	Yield ($/RPM)	LF (%)	Yield ($/RPM)
HF1	85.09	0.1031	88.77	0.0960	86.34	0.1006
HF2	80.38	0.1066	84.86	0.1007	81.15	0.1054
HF3	82.07	0.1038	82.99	0.1022	82.33	0.1033

So HF2 and HF3 overprotect the higher fare classes more than HF1, and IAP1 and IAP2 allow more spiral-down than IAP0. However, the most critical result in this series of experiments is that Airline 1's revenue improvement with FRAT5 "A4" – 3.45% - is greater than that when using FRAT5 "C" – 3.11%.

When shifting from FRAT5 "C" to "A4", Airline 1 assumed more aggressive sell-up in its passengers, and strengthened the booking protections of its higher fare classes. For this particular case, that strategy proved lucrative. So not only can we conclude that HF1-IAP0 constitutes the best Hybrid Forecasting revenue improvement for a given airline using certain FRAT5 estimates, we have determined that there is some relationship between the aggressiveness of the input FRAT5 series and the revenue improvement that airline enjoys. This relationship is further examined in Section 4.3.

4.2.5 Summary of HF1, HF2, HF3 and IAP0, IAP1, IAP2 Tests

In this section, we have tested the nine possible combinations of HF1, HF2, and HF3, and IAP0, IAP1, and IAP2. Each of these methods varies in defining historical bookings as either product-oriented or price-oriented based only on path availability and the state of the next lowest class at the time of booking. And due to this limited booking information, each of these nine methods exhibits some form of bias towards product-oriented or price-oriented demand.

Of the three AP classification methods, IAP1 appears to be the most biased. Based on our understanding of passenger demand (see Section 2.5), we know that the set of all historical bookings in which the next lower class has been closed due to Advance Purchase requirements is comprised exclusively of product-oriented and price-oriented demand. However, we do not know the proportion of these two, and cannot precisely determine it based upon the limited information available.

We intuitively expect the set of AP bookings to contain more price-oriented than product-oriented passengers because selling up to the lowest open fare class is a primary characteristic of price-oriented demand. For this reason, blindly classifying all such bookings as product-oriented (method IAP1) introduces a tremendous amount of bias into the forecasts of future bookings. So it is no surprise that we observe significant spiral-down of revenues with IAP1, and this method was the worst performing IAP classification in nearly every experiment performed here (except for the HF2-IAP1 combination with low demand, in which not enough demand existed to fully initiate the spiral-down effect).

IAP2 seeks to create a medium between the price-oriented focus of IAP0 and the product-oriented focus of IAP1, but does not succeed based upon the experiments performed. Despite the obvious bias in labeling all AP bookings as price-oriented, IAP0 outperforms the others because it minimizes the use of traditional forecasting, thus limiting the spiral-down of revenues.

Regarding the three path availability rules, both the airline rule and the market rule (HF2 and HF3, respectively) appear to be biased because they over-forecast demand in the higher fare classes relative to the path rule (HF1). Doing so consistently leads to overprotection and lower load factors. As described in Section 3.2.3.1, any historical booking able to be classified as product-oriented under HF1 is also product-oriented under HF2 and HF3. Because of the nested definitions of HF1, HF2, and HF3 shown in Figure 45, HF1 will always have the fewest number of historical product-oriented bookings, followed by HF2.

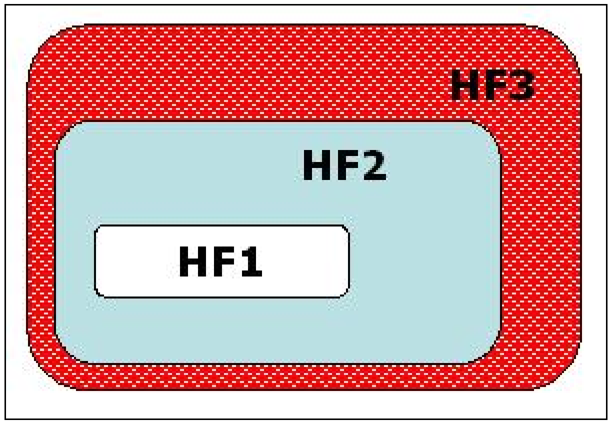

Figure 45: Overlap of HF1, HF2, and HF3 Definitions of Product-oriented Historical Bookings

Therefore, we hypothesize that HF3 leads to higher product-oriented forecasts than HF2, which does the same relative to HF1. Consequently, HF3 and HF2 tend to overprotect the higher fare classes more than HF1 due to their higher forecasts of demand in those classes.

In conclusion, each of these nine methods is biased in that each systematically misidentifies a certain portion of product-oriented and price-oriented bookings. We have simulated each HF method for a DAVN versus DAVN competitive environment of low, baseline, and high demands. We have repeated the experiment in EMSRb versus EMSRb, as well as with a different estimate of passenger WTP. In each of these scenarios, the HF1-IAP0 combination has consistently outperformed all others. For this reason, we believe that Hybrid Forecasting should define its product-oriented and price-oriented historical bookings as follows:

- All bookings in which the next lowest fare class was available on the same path (same flight and airline) are product-oriented.
- All other bookings, including those made when the next lower class has been closed due to AP requirements, are price-oriented.

These definitions constitute our standard methodology for Hybrid Forecasting in PODS (see Figure 28), and were used for all HF experiments in Section 4.3 and Chapter 5 of this thesis.

4.3 Hybrid Forecasting Over a Range of Passenger WTP Estimates

In this section, we examine the relationship between input FRAT5 series and airline network revenue in the context of HF. As discussed in Section 4.2.4, a shift from FRAT5 series "C" to the more aggressive "A4" leads to an increase in Airline 1's revenue of 0.33%, from $1,072,646 to $1,076,200. Relative to AL1's BC1 revenues, FRAT5 series "A4" represents a 3.45% improvement and FRAT5 "C" a 3.11% gain -

so Hybrid Forecasting with either input series represents significant improvements over traditional pick-up forecasting.

4.3.1 Examination of Network Revenues

This situation was repeated sixteen times, with Airline 1 assuming WTPs of FRAT5 series "A9" through "G". As shown in Figure 46, AL1 sees an increase in network revenue with all FRAT5 series except for "G", which results in a loss of nearly 2% from pick-up forecasting. In general, AL1's revenue increases as the input FRAT5 series becomes more aggressive, though the benefit over pick-up forecasting slows around FRAT5 "C". The change in revenue from BC1 significantly jumps from 0.46% with FRAT5 "F" to 3.11% with FRAT5 "C"; yet this benefit only grows to 3.56% with the most aggressive FRAT5 Series "A9", as shown in Table 26. Thus, the revenue benefit from using a WTP estimate more aggressive than FRAT5 "C" is marginal. Furthermore, it is highly unlikely that revenues increase indefinitely with more aggressive sell-up assumptions; however, that inflection point was not yet observed through the range of sixteen FRAT5 series simulated in this thesis.

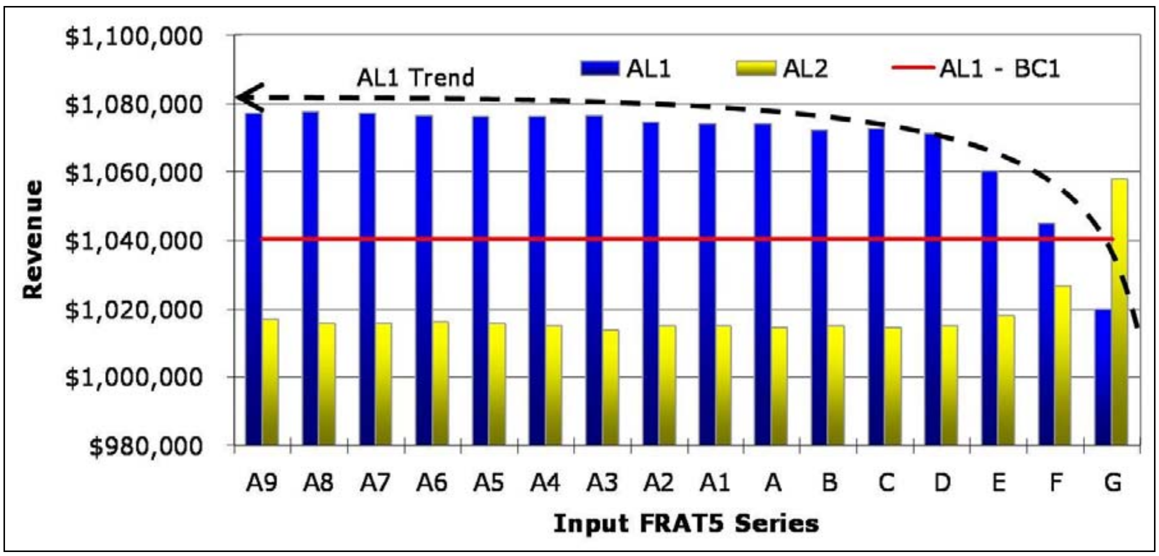

Figure 46: Network Revenue by FRAT5 Series for DAVN w/ HF versus DAVN

Also of note is the inverse relationship shown between AL2's and AL1's network revenues. As AL1's revenues increase with more aggressive FRAT5 series, AL2's decrease – an intuitive relationship in such a competitive network with only two carriers competing on every path.

Table 26: DAVN w/ HF versus DAVN, Airline 1 Revenue Change from BC1 by FRAT5 Series

FRAT5 Series	A9	A8	A7	A6	A5	A4	A3	A2
Δ AL1 Revenue from BC1	3.56%	3.58%	3.55%	3.48%	3.45%	3.45%	3.46%	3.30%

FRAT5 Series	A1	A	B	C	D	E	F	G
Δ AL1 Revenue from BC1	3.27%	3.25%	3.09%	3.11%	2.98%	1.92%	0.46%	-1.98%

4.3.2 Examination of Load Factors and Passenger Yields

As input FRAT5 series become more aggressive, Airline 1's LF decreases significantly from 89.6% to 83.99%, as shown in Table 27. Furthermore, its yield increases from 9.27¢/RPM to nearly 10.5¢/RPM.

Table 27: DAVN w/ HF versus DAVN, Airline 1 Load Factor and Yield by FRAT5 Series

FRAT5 Series	**A9**	**A8**	**A7**	**A6**	**A5**	**A4**	**A3**	**A2**
Load Factor (%)	83.99	84.23	84.37	84.62	84.75	85.09	85.24	85.61
Yield ($/RPM)	0.1046	0.1043	0.1041	0.1037	0.1035	0.1031	0.1029	0.1023

FRAT5 Series	**A1**	**A**	**B**	**C**	**D**	**E**	**F**	**G**
Load Factor (%)	85.77	85.91	86.12	86.2	86.13	86.89	87.86	89.63
Yield ($/RPM)	0.1021	0.1019	0.1015	0.1014	0.1014	0.0995	0.097	0.0927

With a more aggressive FRAT5 series, Airline 1's RM system assumes passengers will engage in more sell-up and increases protections for higher fare classes in turn. Thus, the lower fare classes close sooner and receive fewer bookings – an action which reduces overall load factors, as shown in Figure 47. Of note is the inverse relationship between AL1's and AL2's load factors, which is intuitive and similar to that of their revenue trends.

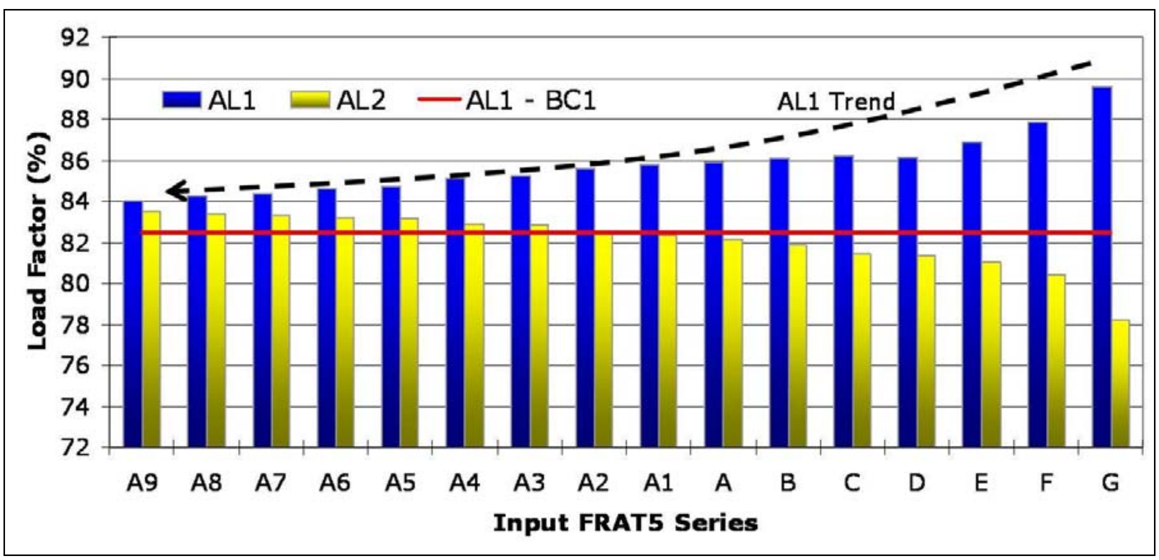

Figure 47: Load Factor by FRAT5 Series for DAVN w/ HF versus DAVN

For all input FRAT5 series, Airline 1's LF exceeds that of BC1, suggesting that HF with any assumption of WTP books more passengers in lower fare classes than pick-up forecasting. Because FRAT5 series "G" loses revenue compared to BC1, its high load factor implies that its WTP estimate is far too low, and the RM system allows excessive low fare class bookings – enough so to dilute revenue and produce a loss.

This least aggressive FRAT5 series also produces the lowest passenger yield for Airline 1. But yields increase as FRAT5 series become more aggressive, as shown in Figure 48. This upward trend in yield is expected with fewer passengers (i.e. lower load factors) and increasing network revenues as HF assumes more aggressive sell-up.

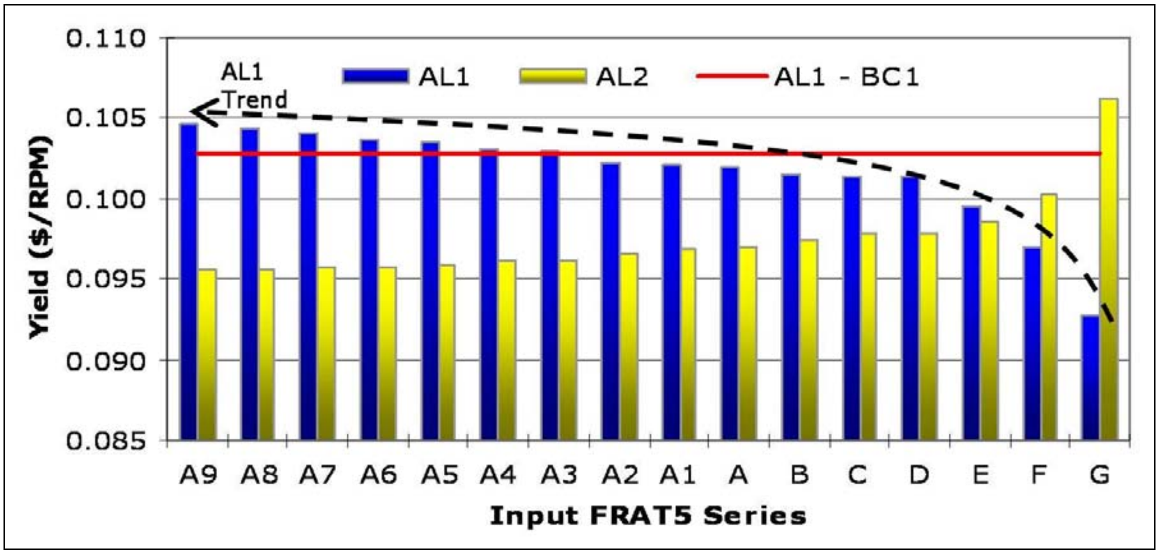

Figure 48: Yield by FRAT5 Series for DAVN w/ HF versus DAVN

4.3.3 Summary of DAVN w/ HF versus DAVN

In closing, the use of HF nearly always increases network revenues over traditional pick-up forecasting, except in cases where the assumption of passenger sell-up is so conservative that revenues are significantly diluted. As inputs of WTP increase, the airline's RM system becomes more protective of its higher fare classes. Doing so limits bookings in lower fare classes and drives load factors downward. Consequently, the average revenue generated per passenger increases due to the reduced number of bookings. And as shown in this section, network revenue continues to improve with more aggressive WTP estimates, though the incremental benefit declines near the level of FRAT5 series "C", and that an Airline using HF may see an increase in revenue approaching 3% over pick-up forecasting.

4.4 Chapter Summary: Simulations in PODS to Define HF

This chapter first defined HF and then determined the revenue improvements possible relative to standard pick-up forecasting. We began in Section 4.1 by presenting the simulated network in PODS – Network D-6 – as well as defining the semi-restricted fare structure used throughout this thesis.

Then in Section 4.2, we experimented with the nine historical booking definitions for use in HF, as previously described in Section 3.2.3.1 and 3.2.3.2. Specifically, we tested each of the nine scenarios in environments of baseline, low, and high demand within PODS, as well as with the leg-based EMSRb optimizer in lieu of DAVN, and with a more aggressive estimate of passenger WTP. The results of our experiments indicate that HF1-IAP0 outperforms the other 8 definitions for two reasons:

1. HF1 is characterized by less overprotection than HF2 and HF3.
2. IAP0 better guards against spiral-down of revenues than IAP1 and IAP2.

Using this HF1-IAP0 combination to define product-oriented and price-oriented historical bookings, we then simulated HF in Section 4.3 to demonstrate the sensitivity of HF to an estimate of passenger WTP. Based on our experiments, it appears that an airline using Hybrid Forecasting can experience an increase in network revenue of approximately 3% over traditional forecasting methods.

In Chapter 5, we will present further experiments using PODS in which we can improve the performance of Hybrid Forecasting in terms of further increasing network revenues. Specifically, we will demonstrate the value of varying the input WTP estimate by relative path quality – a technique referred to as Path Categorization. Furthermore, we simulate another supplemental technique known as Fare Adjustment, which proactively accounts for passenger sell-up in the RM system's seat optimizer (instead of the forecasting component, like HF).

5 Improving Hybrid Forecasting using Path Categories and Fare Adjustment

In Chapter 4, we showed that the HF1-IAP0 combination often outperforms the other eight categorizations of product-oriented and price-oriented bookings as described in Sections 3.2.3.1 and 3.2.3.2, respectively. We then demonstrated the sensitivity of an airline's network revenues to its estimated passenger Willingness-to-Pay (WTP) inputs for Hybrid Forecasting (HF) with less restricted fare structures.

Having demonstrated the potential value of HF to an airline in an environment with semi-restricted fare structures, our goal in this chapter is to expand our experiments to test additional improvement in Hybrid Forecasting in terms of increasing network revenue derived from HF alone. More specifically, we present and analyze simulations of an airline using HF supplemented with two techniques: Path Categorization (PCAT, as described in Section 3.2.5) and Fare Adjustment (FA, as described in Section 3.2.4).

As in Chapter 4, we use the Passenger Origin-Destination Simulator (PODS) to simulate the performance of HF, PCAT, and FA within an airline Revenue Management (RM) system in semi-restricted fare structures. Again, we refer to experiments in terms of "the RM system used for Airline 1 (AL1) versus that used for Airline 2 (AL2)" in semi-restricted Network D-6; both carriers employ the Displacement Adjusted Virtual Nesting (DAVN) seat allocation optimizer, as described in Section 3.1.3.3. Also, AL2 always uses standard pick-up forecasting, while the forecasting method used by AL1 varies in each experiment.

We begin by we analyze the impact of Hybrid Forecasting with PCAT in terms of path quality in Section 5.1; next, we examine the effect of FA on HF in Section 5.2; and finally, we study the combination of PCAT and FA together, and their combined impact on HF in Section 5.3.

5.1 Hybrid Forecasting using Path Categories to account for relative path quality

In this section, we test the impact of an airline accounting for different passenger sell-up behavior in markets with non-stop service. As mentioned in Section 1.3.2 and discussed in Section 3.2.5, we intuitively expect that an airline offering superior service to its competition in terms of path quality (non-stop versus connecting) can extract more revenue because of passenger preference for non-stop service. That is, passenger WTP should be higher for non-stop service compared to connecting.

We refer to the set of Origin-Destination (OD) markets where a given airline dominates in terms of path quality as "PCAT1". Similarly, the OD markets where all competitors offer equal path quality (all non-stop or all connecting service) are "PCAT2", and the OD markets in which a given airline is dominated by a competitor in terms of path quality (connecting versus non-stop) are "PCAT3". And in Network D-6, each airline in each bank has 40 PCAT1 paths out of 482, as discussed in Section 4.1.1.

Due to the preference for non-stop service, we expect a greater occurrence of sell-up in PCAT1 if the lower fare classes are proactively closed due to a lack of competing non-stop flights. In terms of our simulated PODS environment, we test this hypothesis by having AL1 (using HF) selectively use more aggressive FRAT5 inputs for PCAT1 than for PCAT2 and PCAT3.

Regarding the PCAT experimentation nomenclature in this thesis, we describe each combination as "PCAT1 FRAT5 – PCAT2 FRAT5 – PCAT3 FRAT5". For example, input FRAT5 series combination "A4-C-C" uses FRAT5 series "A4" for PCAT1 and "C" for PCAT2 and PCAT3.

5.1.1 Path Categorization – Network Revenue

To isolate the effects of PCAT, we return to our baseline input of FRAT5 series "C" (see Section 4.1.3). Specifically, we fix FRAT5 "C" for PCAT2 and PCAT3, and vary the FRAT5 series used for PCAT1 to gauge the sensitivity of AL1's network revenue to more aggressive input FRAT5.

The Base Case (BC) for this scenario is "C-C-C", where AL1 uses FRAT5 series "C" for every path irrespective of path quality. This particular BC was presented previously in Section 4.3 (where no PCAT was used), and is summarized in Table 28.

Table 28: Base Case "C-C-C" Results – DAVN w/ HF and PCAT versus DAVN

Airline	Revenue	Load Factor (%)	Yield ($/RPM)
AL1	$1,072,646	86.20	0.1014
AL2	$1,014,617	81.42	0.0978

In this Base Case, Airline 1 gets $1,072,646 in network revenue – 3.11% more than in BC1 (see Chapter 4) – because of its use of HF instead of pick-up forecasting. Also, AL1's load factor (LF) is higher than in BC1 because the use of HF leads to more bookings in lower fare classes, as discussed in Section 4.3. For the experiments in this section, we are concerned with the possible revenue improvements for AL1 when using more aggressive FRAT5 series for PCAT1.

This situation was repeated thirteen times, with Airline 1 assuming WTP estimates for PCAT1 of FRAT5 series "A8" through "E". As shown in Figure 49, AL1 generally sees an increase in network revenue as its PCAT1 FRAT5 series becomes more aggressive – consistent with our intuitive expectation.

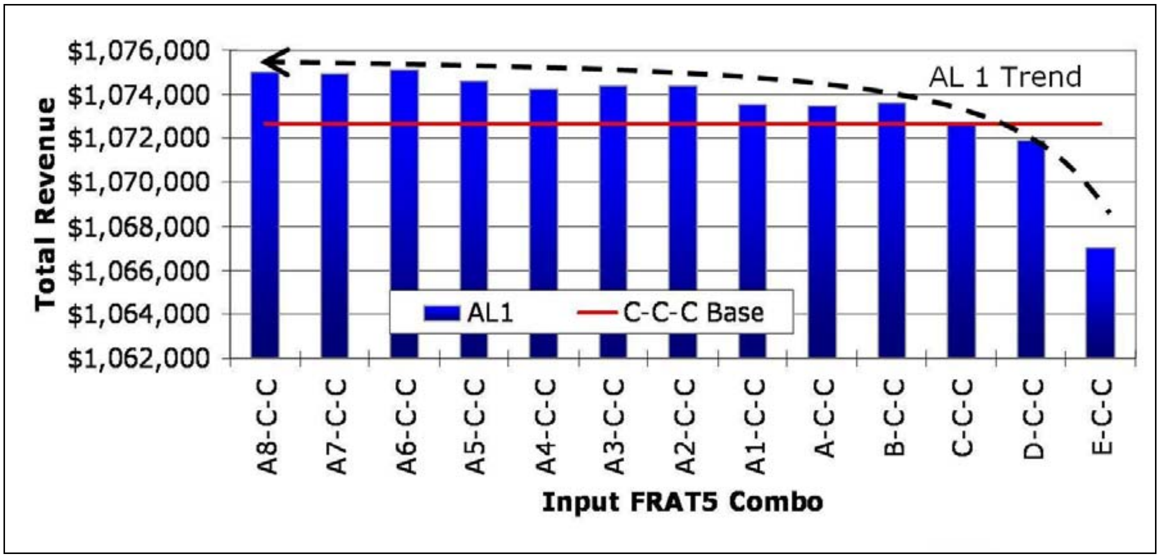

Figure 49: AL1 Revenue by FRAT5 Series Combination, DAVN w/ HF and PCAT versus DAVN, Using "C" for PCAT2 and PCAT3

This revenue increase due to PCAT is relatively small – less than 0.25% over the "C-C-C" Base Case, as shown in Table 29 – compared to the 3.11% revenue increase due to the use of HF alone. The smaller magnitude of this revenue change can be explained by the small number of OD paths in PCAT1 for AL1 in Network D-6. Note that less than 9% of AL1's paths exhibit superior path quality to AL2's competing service.

Table 29: DAVN w/ HF and PCAT versus DAVN, Airline 1 Revenue Change by PCAT1 FRAT5 Series, Using "C" for PCAT2 and PCAT3

PCAT1 FRAT5 Series	A8	A7	A6	A5	A4	A3	A2
Δ AL1 Revenue from "C-C-C"	0.22%	0.21%	0.23%	0.18%	0.15%	0.16%	0.16%
Δ AL1 Revenue from BC1	3.34%	3.33%	3.35%	3.30%	3.26%	3.28%	3.28%

PCAT1 FRAT5 Series	A1	A	B	C	D	E
Δ AL1 Revenue from "C-C-C"	0.09%	0.07%	0.09%	-	-0.07%	-0.52%
Δ AL1 Revenue from BC1	3.20%	3.19%	3.21%	3.11%	3.04%	2.57%

So even though only 40 of AL1's 482 paths saw a change in input passenger WTP estimates, the airline still experienced an increase in total network revenue of up to 0.23% (with the best case input of FRAT5 "A6"). This finding indicates that supplementing Hybrid Forecasting with Path Categorization can be an effective way to improve an airline's revenue with only minor changes to the RM system, though the small magnitude of the gain leaves us unable to conclude the revenue gain is due exclusively to passenger sell-up in these PCAT1 paths.

5.1.2 Path Categorization – Yield and Load Factor

Because of this increase in revenue, we intuitively expect that PCAT captures sell-up in these 40 OD markets and thereby increases the average fare paid per passenger.

That is, we expect overall passenger yield to increase due to sell-up in the 40 PCAT1 markets. We further expect no increase in Airline 1's load factor because sell-up does not attract additional passengers, but rather extracts additional revenue from existing passengers. If anything, we expect a slight decrease in LF due to the use of PCAT, as the more aggressive booking limits in the PCAT OD markets may dissuade existing passengers from booking. However, our intuitive expectations for both yield and load factor appear to be incorrect based on our experiments.

In contrast to expected enhancement in passenger yield, we observe no discernable trend, as shown in Figure 50. For all PCAT FRAT5 series "A8" through "C", the yield only varies between 10.13 ¢/RPM and 10.15 ¢/RPM – not a significant departure from the 10.14 ¢/RPM yield of the "C-C-C" base.

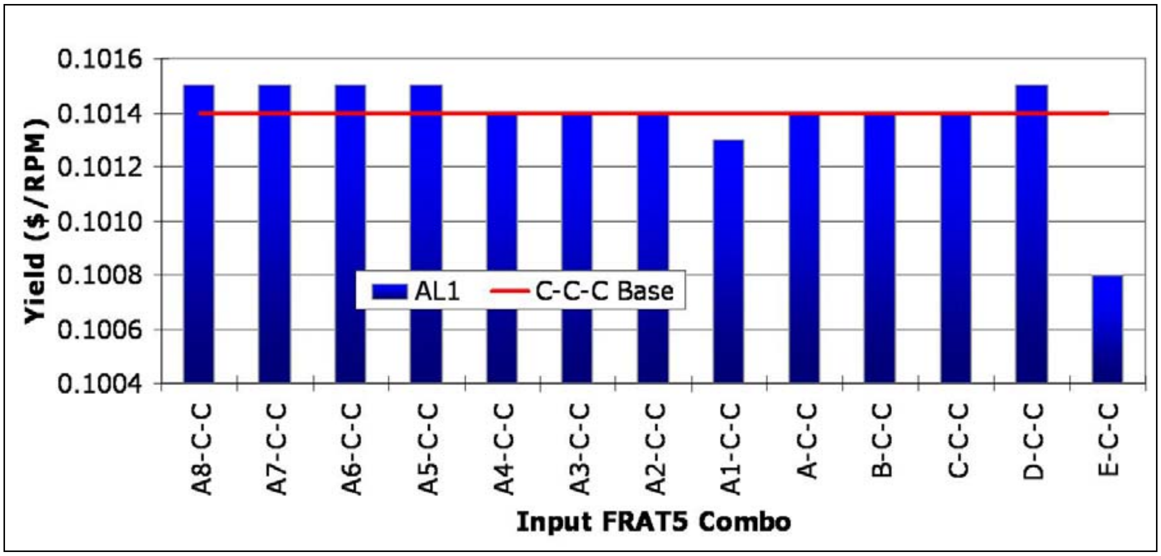

Figure 50: AL1 Yield by FRAT5 Series Combination, DAVN w/ HF and PCAT versus DAVN, Using "C" for PCAT2 and PCAT3

And instead of Airline 1's load factor remaining stable, or even slightly decreasing with more aggressive WTP estimates for PCAT1, we observe a slight LF increase over the "C-C-C" Base Case for all FRAT5 series more aggressive than "C" as shown in Figure 51. This finding is counter-intuitive because the use of more aggressive FRAT5s on PCAT1 paths should result in stronger booking class protection levels and more denied bookings (thus, lower load factors).

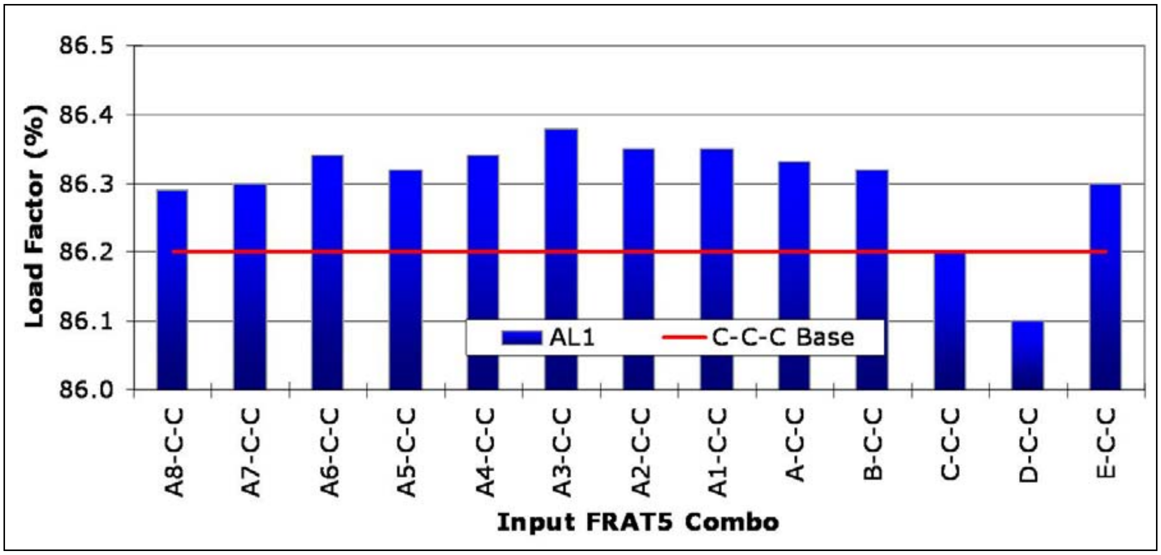

Figure 51: AL1 LF by FRAT5 Series Combination, DAVN w/ HF and PCAT versus DAVN, Using "C" for PCAT2 and PCAT3

To explain these results, we examine a representative case of Hybrid Forecasting with PCAT and observe a complex interaction between Airline 1 and Airline 2. With the input FRAT5 combination "A2-C-C", Airline 1's network revenues exceed those of the "C-C-C" Base Case by 0.16%, or $1,704, as shown in Table 29. In Table 30, we examine the sources of this revenue gain using Path Categories defined by relative path quality.

When using PCAT, Airline 1 loses $11,686 from "1st Choice" passengers who are denied booking on their preferred OD path and fare class. for an example. But note that over $10,000 of this loss occurs in the lowest two fare classes – an indication that the use of PCAT protects higher fare classes by originally denying bookings in the lower ones. Furthermore, note that AL1 "recaptures" $8,967 more from these passengers by booking them on a different fare class or path when it uses PCAT than otherwise.

Table 30: Changes in Airline Revenue by Fare Class, DAVN w/ HF, from "C-C-C" Base to "A2-C-C"

Fare Class	Total	1st Choice	Sell-up	AL1 Recapture	AL2 Spill-in
1	$1,722	$350	$653	$205	$515
2	$2,546	$352	$1,227	$615	$351
3	-$1,328	-$2,067	-$77	$1,049	-$231
4	$1,936	$150	$148	$1,284	$354
5	-$7,583	-$8,123	-$1,721	$1,676	$584
6	$4,411	-$2,348	$0	$4,138	$2,622
Σ	$1,704	-$11,686	$230	$8,967	$4,195

Also, note that AL1 captures $1,880 more from sell-up to Fare Classes 1 and 2 when using "A2" in PCAT1 – an indication that PCAT leads to sell-up in the higher fare classes as originally hypothesized. But the changes in 1st Choice and Recapture revenue for Airline 1, as well as the additional $4,195 captured from AL2's denied passengers indicates that the use of FA leads to gains in network revenue due to network effects beyond capturing sell-up in PODS Network D-6. While we can observe the resultant change in revenue due to these network effects, it is clear that the use of Path Categories has an indeterminate effect on load factors and yields. However, we can conclude that this effect is somewhat muted due to AL1's limited number of PCAT1 paths in Network D-6.

5.1.3 Conclusions about Path Categorization

Based upon the trends in Airline 1's revenue, load factor, and yield over a range of PCAT1 WTP estimates, our simulations support the following two conclusions:

1. An airline practicing PCAT by relative path quality can experience a significant positive impact on its total network revenues to due the higher WTP in a relatively small number of OD markets. As shown in Table 29, this increase in revenue can approach 0.25% over Hybrid Forecasting alone.
2. Despite possibly increasing revenues via sell-up in PCAT1 markets, the revenue benefits of PCAT appear to be due to network effects when one airline assumes higher WTP inputs in its dominant markets. Due to these network effects in our experiments, we observe that load factors increase slightly due recapturing of denied bookings from both AL1 and AL2.

5.2 Hybrid Forecasting using Fare Adjustment to account for passenger sell-up

In this section we turn our attention to Fare Adjustment as another supplemental tool for improving the performance of Hybrid Forecasting. Specifically, we demonstrate through simulation the incremental effect that FA has on HF (and vice versa) in terms of revenue for Airline 1. To put these changes in context, we first demonstrate the effects that FA alone has on an airline in Section 5.2.1, and then we show the cumulative result of both HF and FA together in Section 5.2.2.

As first introduced in Section 1.3.1, and then described in Section 2.4.2 and Section 3.2.4, Fare Adjustment is another method of accounting for passenger sell-up in less restricted fare structures. While Q-forecasting – the price-oriented component of Hybrid Forecasting – was developed to incorporate sell-up into the passenger demand forecasting step of RM, as discussed in Section 2.4.1 and Section 3.2.2, Fare Adjustment takes place within the RM system's seat allocation optimizer. For use within a virtual nesting environment, FA artificially lowers the fare of the lower classes to shift them into lower virtual nests. Doing so closes these classes earlier and theoretically encourages sell-up into the higher fare classes.

Because this passenger sell-up behavior is governed by a WTP estimate – or a FRAT5 value in PODS – the aggressiveness of an airline's Fare Adjustment methods must be dependent upon that airline's estimate of passenger WTP. In PODS, we continue to use the FRAT5 series first described in Section 3.2.2.1 to estimate WTP; but for FA, we apply a FRAT5 Scaling Factor, as described in Section 3.2.4.2. So an airline in PODS practicing Fare Adjustment must now estimate two values: the overall FRAT5 describing passenger WTP, and the FRAT5 Scaling Factor used exclusively for Fare Adjustment.

As an example of FA used in our simulations, consider the cases of FRAT5 series "C" and "A4", both with a Scaling Factor of 0.25. The examples presented below are for Airline 1 in the OD market from City 12 to Hub 1, as originally described in Section 4.1.1's discussion of Network D-6. The adjusted fares, *fare'*, to be used in determining DAVN booking limits for each class and each Time Frame are shown in Figure 52 for FRAT5 "C" and in Figure 53 for FRAT5 "A4".

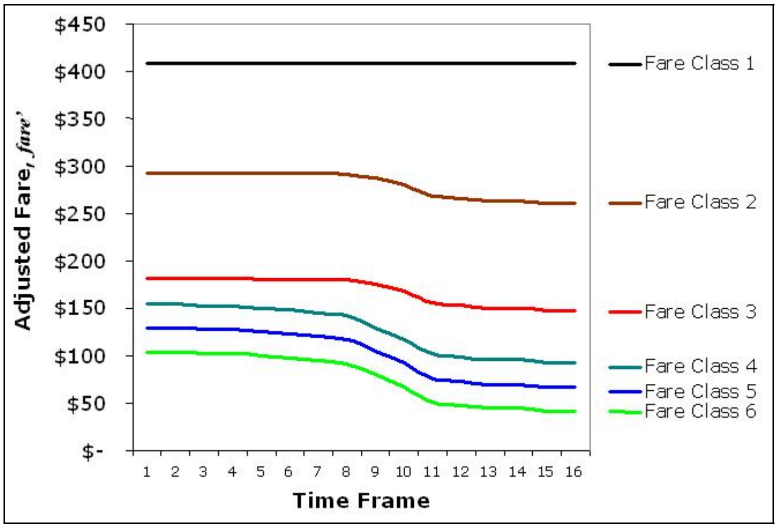

Figure 52: Airline 1 Adjusted Fares for City 12-H1 OD Market, FRAT5 "C", FRAT5 Scaling Factor of 0.25

The characteristic "S-shape" of our FRAT5 series is mirrored for the adjusted fares in both Figure 52 and Figure 53. Also, note that no Fare Adjustment is ever performed on the highest fare class (Fare Class 1) because no sell-up to an even higher class is possible. The actual fares paid by passengers, as well as the final adjusted fares used by the DAVN optimizer in Time Frame 16 are shown in Table 31.

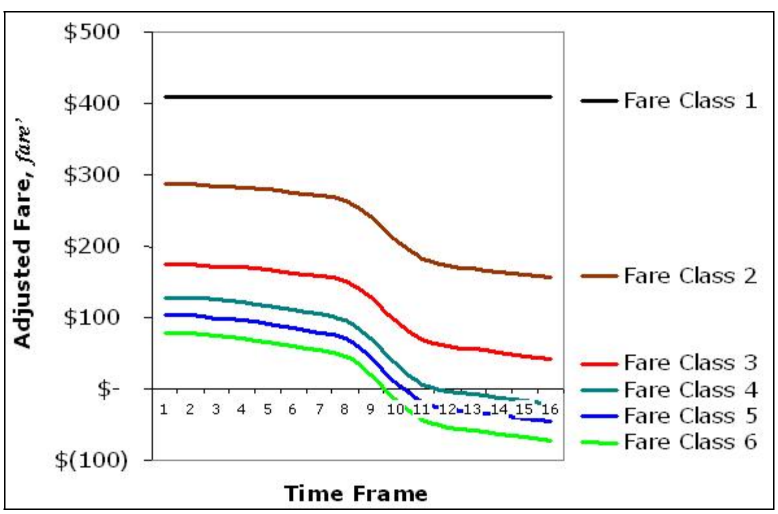

Figure 53: Airline 1 Adjusted Fares for City 12-H1 OD Market, FRAT5 "A4", FRAT5 Scaling Factor of 0.25

Table 31: Adjusted Fares in Time Frame 16 for Airline 1, City 12-H1 OD Market, FRAT5 Scaling Factor of 0.25

Fare Class	1	2	3	4	5	6
Actual Fares	$408.90	$292.40	$181.00	$154.70	$129.40	$104.10
FRAT5 "C" - TF 16	$408.90	$260.50	$147.70	$91.90	$66.50	$41.20
FRAT5 "A4" - TF 16	$408.90	$155.40	$41.90	-$21.30	-$46.70	-$72.00

In comparing these two FRAT5 series, use of the more aggressive one ("A4") leads to more aggressive Fare Adjustment in which the adjusted fares experience larger decreases (to the point of negative adjusted fares) than with the less aggressive series ("C"). So if Airline 1 assumes FRAT5 series "A4" for passenger WTP, its FA will shift fare classes into lower virtual nests sooner than if Airline 1 had used "C". In doing so, these lower fare classes will close earlier in the booking process, and the airline will (it hopes) capture sell-up to the higher fare classes.

The experiments described in Section 5.2.1 simulate FA alone, and are performed over a range FRAT5 Scaling Factors for FRAT5 series "C" and "A4". These experiments are then repeated in Section 5.2.2, but with HF added to FA. To judge the impact of FA, each experiment must be compared to an appropriate Base Case without Fare Adjustment, or the "No FA Base". Of course, this "No FA Base" is different for each experiment, and will be described in turn.

5.2.1 DAVN w/ FA versus DAVN – No Hybrid Forecasting

Because HF is not used in this section, the appropriate "No FA Base" to measure the impact of FA alone is that of DAVN versus DAVN, which has been already presented in Chapter 4 as Base Case 1 (see Table 9). Airline 1's network revenue in BC1 is $1,040,277, which beats AL2 by $9,727 due to asymmetries in Network D-6. In terms of Fare Adjustment, BC1 is equivalent to Airline 1 practicing FA with a FRAT5 Scaling Factor of 0.0 (the choice of FRAT5 Series is irrelevant).

5.2.1.1 Using FRAT5 Series "C" for Fare Adjustment

For this set of experiments, we again return to our baseline FRAT5 series "C", but now apply FA with ten gradually increasing Scaling Factors, starting with 0.05 and incrementing by 0.05 to a maximum of 0.50. As this Scaling Factor grows, our FA becomes more aggressive, and the DAVN optimizer becomes more proactive in closing lower fare classes.

Intuitively, we expect Airline 1's network revenue to increase now that the FA process accounts for passenger sell-up in the semi-restricted fare structure. But as illustrated in Figure 54, the simulations show counter-intuitive results. The biggest beneficiary from AL1's FA is actually its competitor, whose revenues increase tremendously as the FA becomes more aggressive.

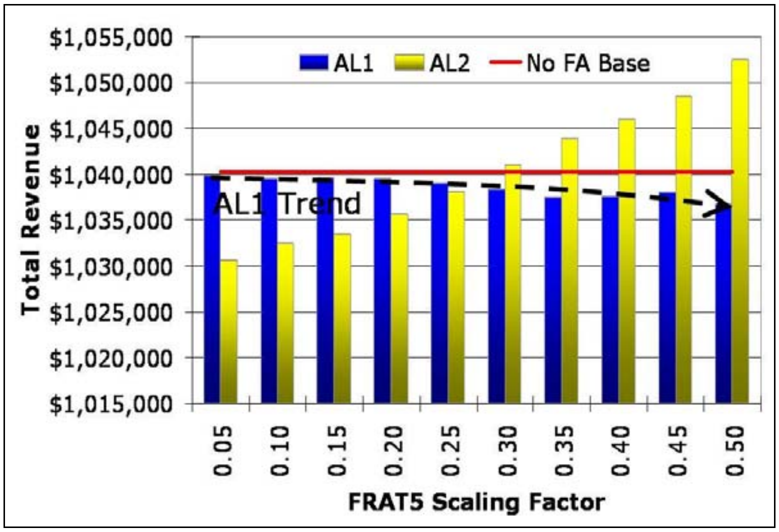

Figure 54: Network Revenue by FRAT5 Scaling Factor for DAVN w/ FA versus DAVN, FRAT5 "C"

The most startling result of these experiments is that the introduction of FA – a more sophisticated RM development – actually hurts AL1. Note that the type of FA applied here is "universal" in that the same input FRAT5 series is used for all paths, in contrast to the relative path quality experiments performed in Section 5.1. Airline 1's revenues slightly decrease from the "No FA Base" under all forms of universal FA. And as shown in Table 32, the revenue loss becomes magnified as more aggressive FA is used, growing (or actually declining) from 0.03% with the smallest Scaling Factor to 0.34% with a Scaling Factor of 0.50.

Table 32: DAVN w/ FA versus DAVN, Airline 1 Change in Network Revenue from BC1 by FRAT5 Scaling Factor, FRAT5 "C"

FRAT5 Scaling Factor	**0.05**	**0.10**	**0.15**	**0.20**	**0.25**
Δ AL1 Revenue from BC1	-0.03%	-0.07%	-0.07%	-0.07%	-0.13%
FRAT5 Scaling Factor	**0.30**	**0.35**	**0.40**	**0.45**	**0.50**
Δ AL1 Revenue from BC1	-0.18%	-0.27%	-0.26%	-0.22%	-0.34%

To explain these counter-intuitive trends, we examine both airlines' load factors and yields in these ten FA experiments. As shown in Figure 55, AL1's load factor decreases almost linearly (and significantly) as FA becomes more aggressive. Due to this heavy reduction in bookings (and only a slight decrease in revenues), AL1's yield rises sharply with the Scaling Factor used for FA.

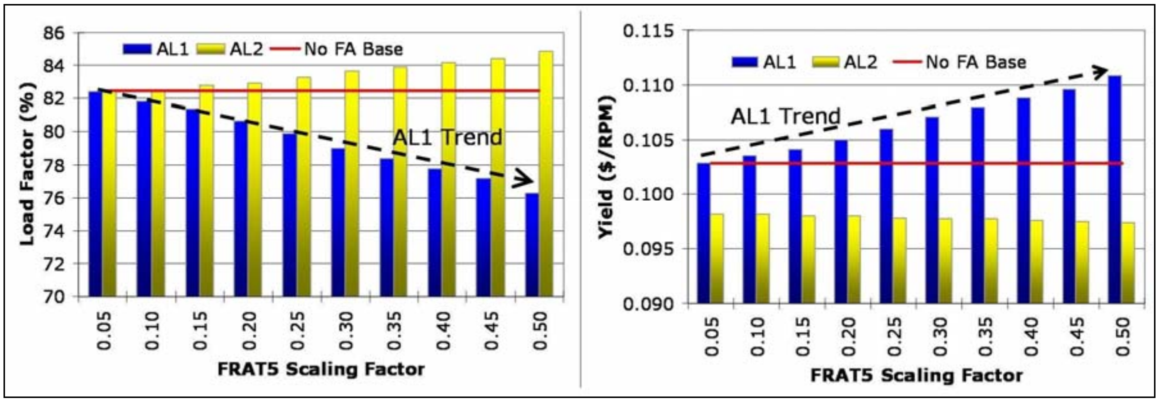

Figure 55: Load Factor and Yield by FRAT5 Scaling Factor for DAVN w/ FA versus DAVN, FRAT5 "C"

This combination of AL1's revenues, load factors, and yields indicates that FA causes the DAVN seat optimizer to overprotect the higher fare classes. By adjusting fares too quickly, FA causes the optimizer to close lower classes too early in hopes of attracting high fare bookings which do not materialize. The loss of network revenue due to FA only grows with use of more aggressive Scaling Factors. So as the FA becomes more aggressive, the lower fare classes close earlier, and Airline 1's load factor plummets due to the loss of so many low revenue bookings. Consequently, its yield increases. But those extra bookings in the higher fare classes – passengers the RM system sold up – cannot offset the revenue lost due to overprotection.

Conversely, AL2 increases its load factors by capturing some of its competitor's spilled passengers. And because its own RM system manages to keep passenger yield from significantly decreasing with the additional passengers, Airline 2 experiences the remarkable increases in revenue observed in Figure 54, and even manages to pass its competitor when AL1 uses a Scaling Factor of 0.30 or higher.

5.2.1.2 Using FRAT5 Series "A4" for Fare Adjustment

After observing Airline 1's slight (and unexpected) revenue loss due to universal FA with FRAT5 "C", we repeat the experiments to test the airline's sensitivity to FA with a different input WTP estimate. And with AL1 using the more aggressive FRAT5 series "A4", the loss of network revenue from the "No FA Base" becomes even more pronounced. As shown in Figure 56, Airline 2 again experiences a surge in revenue due to Airline 1's use of FA.

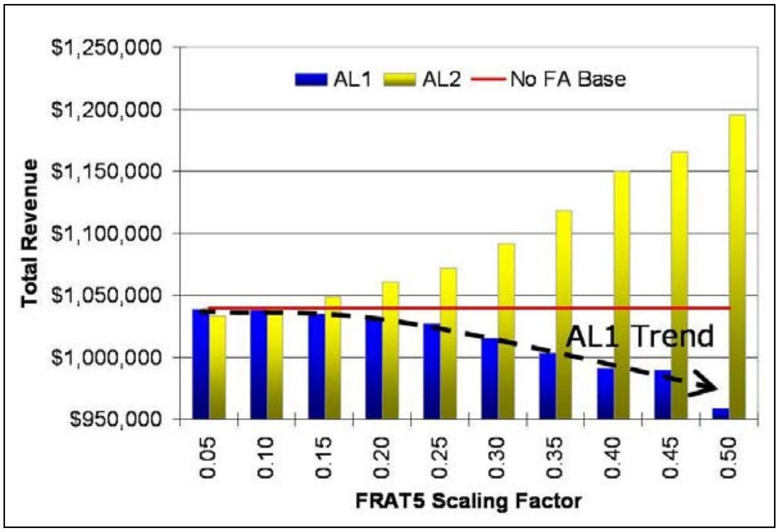

Figure 56: Network Revenue by FRAT5 Scaling Factor for DAVN w/ FA versus DAVN, FRAT5 "A4"

At a Scaling Factor of 0.50, AL1 sees a revenue loss of nearly 8% (see Table 33) from the case of no FA – much higher than the loss of 0.34% at the same scaling factor with FRAT5 "C". So as demonstrated in the case of FRAT5 "C", more aggressive FA – this time by way of a more aggressive FRAT5 – leads to a heavier loss of revenue for AL1 (as well as a greater increase for AL2).

Table 33: DAVN w/ FA versus DAVN, Airline 1 Change in Network Revenue from BC1 by FRAT5 Scaling Factor, FRAT5 "A4"

FRAT5 Scaling Factor	**0.05**	**0.10**	**0.15**	**0.20**	**0.25**
Δ AL1 Revenue from BC1	-0.13%	-0.18%	-0.49%	-0.85%	-1.25%

FRAT5 Scaling Factor	**0.30**	**0.35**	**0.40**	**0.45**	**0.50**
Δ AL1 Revenue from BC1	-2.42%	-3.50%	-4.73%	-4.88%	-7.78%

In terms of load factor and yield, we notice the same trends with FRAT5 "A4" as we did with "C" – only intensified, as shown in Figure 57. In this case, AL1's LF falls precipitously – approaching 50% with the highest Scaling Factor of 0.50. And AL1's yield skyrockets, approaching 16 ¢/RPM at the same high Scaling Factor.

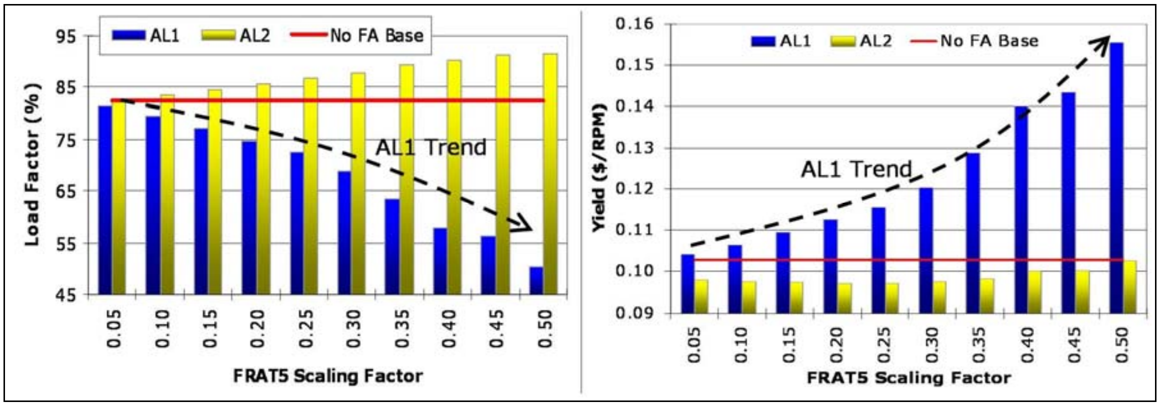

Figure 57: Load Factor and Yield by FRAT5 Scaling Factor for DAVN w/ FA versus DAVN, FRAT5 "A4"

Of course, AL2's load factors creep upward accordingly as it captures AL1's spilled booking requests. Interestingly, AL2's yields actually begin to rise too as the FA becomes extremely aggressive (Scaling Factors above 0.30), demonstrating that AL1 is rejecting so many potential bookings, that AL2 can actually afford to reject some of those same spilled (albeit low fare) bookings as well.

5.2.1.3 Conclusions about FA (without HF)

In summary, it appears that Fare Adjustment on its own is an unappealing method for capturing sell-up in semi-restricted fare structures, despite its expected benefits. In the simulations performed here, FA actually leads to revenue losses for Airline 1, as shown in Figure 58. And as the FA becomes more aggressive – both in terms of the Scaling Factor and the FRAT5 Series used to estimate passenger WTP – those losses become more severe.

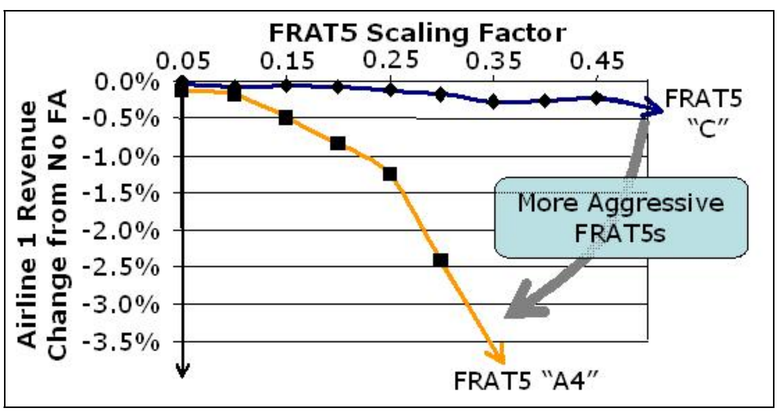

Figure 58: Change in Airline 1 Network Revenue by FRAT5 Scaling Factor, from DAVN (Base Case 1) to DAVN w/ FA versus DAVN

Theoretically designed to encourage sell-up by closing lower fare classes earlier, Fare Adjustment (on its own) appears to be too aggressive in doing so. As shown in our experiments, FA leads the DAVN optimizer to close these lower fare classes too soon, decreasing load factor and increasing yield. The expected bookings in higher fare classes fail to materialize – a case of revenue loss due to overprotection, as originally shown in Figure 2.

5.2.2 <u>DAVN with HF and FA together versus DAVN</u>

Despite the ineffectiveness of universal FA alone, we now examine the use of FA and Hybrid Forecasting in conjunction with an aim of improving network revenues. As with the previous section, we first present HF and FA using two FRAT5 series: "C" and "A4". Doing so allows us a direct comparison of FA with and without HF for both of these series. We also perform the complementary comparison of HF with and without FA, which allows us to measure the benefit of FA as a tool to improve the performance Hybrid Forecasting in semi-restricted fare structures – one of the goals of this thesis.

5.2.2.1 Using FRAT5 Series "C" for Hybrid Forecasting with FA

We begin by experimenting exclusively with the moderate FRAT5 series "C" as Airline 1's estimate of passenger WTP. In this case, the appropriate "No FA Base" is that of DAVN w/ HF versus DAVN where Airline 1 employs FRAT5 "C". Note that this scenario was presented previously in Section 4.3 and Section 5.1, and is summarized in Table 28. With only HF, Airline 1 sees network revenues of $1,072,646 – significantly higher than Airline 2's total of $1,014,617 (without the benefit of HF).

As in Section 5.2.1.1, we apply FA ten times with Scaling Factors of 0.05 through 0.50. The aggressiveness of the HF does not change in any of these experiments; however the FA becomes more aggressive with increasing Scaling Factors.

Intuitively, we expect revenues to increase over the "No FA Base" with the introduction of Fare Adjustment to Airline 1's RM system. As discussed in Chapter 4, HF is an effective way to capture passenger sell-up in semi-restricted fare structures. Fare Adjustment manages sell-up somewhat independently of HF because it operates within the RM system's seat optimizer component (as opposed to the forecaster). Because we expect that HF gives the DAVN optimizer a better forecast of demand in each fare class (at least in semi-restricted fare structures) than traditional pick-up forecasting, we expect an improvement over HF alone. That is, we intuitively believe that the use of HF will help guard against the overprotection of FA alone, as described in Section 5.2.1, resulting in an increase in network revenues.

As shown in Figure 59, our expectations are realized for AL1 – at least through a maximum Scaling Factor of 0.50. Using FA and HF together leads Airline 1 to network revenue gains for each FA experiment compared to the appropriate "No FA Base".

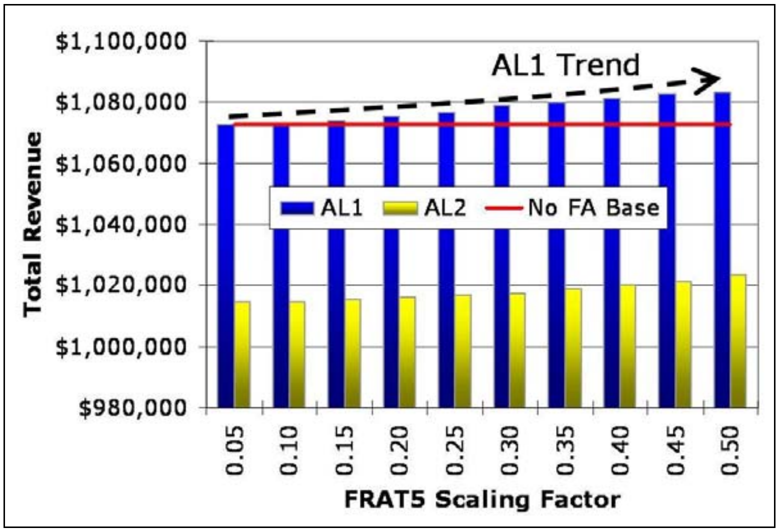

Figure 59: Network Revenue by FRAT5 Scaling Factor for DAVN w/ HF and FA versus DAVN, FRAT5 "C"

As shown in Table 34, AL1's potential revenue increase grows as the Fare Adjustment becomes more aggressive – up to 1% over Hybrid Forecasting alone using the maximum Scaling Factor tested of 0.50. Indeed, this gain in network revenue using FA to supplement HF is a departure from the unexpected losses of FA by itself demonstrated previously.

Table 34: DAVN w/ HF and FA versus DAVN, Airline 1 Change in Network Revenue from DAVN w/ HF by FRAT5 Scaling Factor, FRAT5 "C"

FRAT5 Scaling Factor	**0.05**	**0.10**	**0.15**	**0.20**	**0.25**
Δ AL1 Revenue from HF Alone	0.00%	0.05%	0.13%	0.25%	0.39%
FRAT5 Scaling Factor	**0.30**	**0.35**	**0.40**	**0.45**	**0.50**
Δ AL1 Revenue from HF Alone	0.59%	0.71%	0.82%	0.95%	1.00%

Also of interest in Figure 59 is that Airline 2 sees revenue gains due to its competitor's use of FA. Typically, we expect that a gain by one airline results in a loss for its opposite number. To understand why we observe revenue gains for each airline, we again turn to their load factors and yields, as shown in Figure 60.

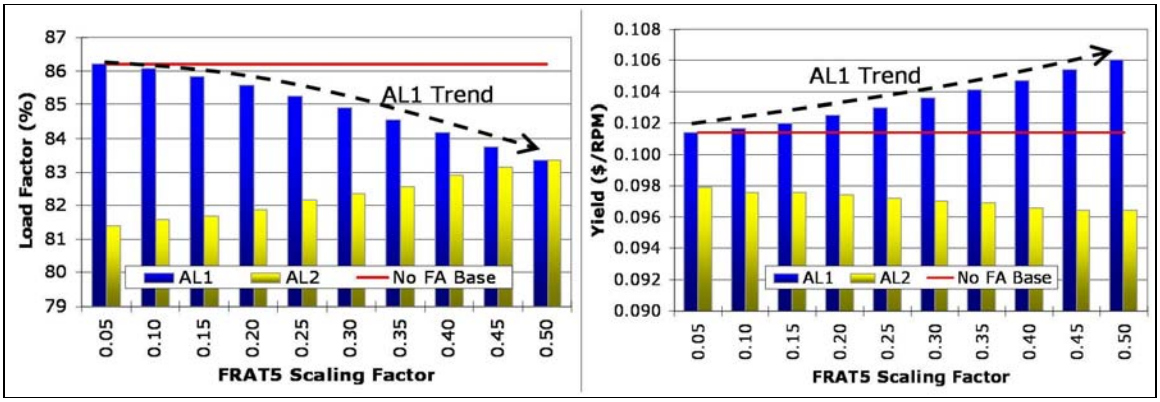

Figure 60: Load Factor and Yield by FRAT5 Scaling Factor for DAVN w/ HF and FA versus DAVN, FRAT5 "C"

As was the case with FA alone, Airline 1's load factor decreases as FA becomes more aggressive. But the use of Hybrid Forecasting lessens the severity of these LF decreases. For example, FA alone saw a decrease in LF of 7.5% - from 82.46% in the "No FA Base" to 76.25% with Scaling Factor 0.50. With HF and FA simultaneously used by AL1's RM system, LF only decreased by 3.3% - from 86.2% in its respective "No FA Base" to 83.35% with the same Scaling Factor. Airline 1's combination of reduced loads and passenger sell-up due to FA leads to higher yields as well. And based on the overall increase in revenue, it appears that Airline 1's use of FA and HF together is not characterized by nearly as much overprotection as that seen with FA alone.

Regarding Airline 2, we observe increasing load factors due to AL1's FA, as well as slightly decreasing yields. But because its own network revenue slightly grows as AL1's FA becomes more aggressive, we conclude that Airline 2's revenue in the "No FA Base" was characterized by some overprotection as well. With Airline 1 now spilling more passengers due to FA – especially as FA becomes more aggressive – AL2 captures enough of this spill to increase its revenues as well. This is an example of a situation where both competitors improve their situations due to the strategic actions of only one.

5.2.2.2 Using FRAT5 Series "A4" for Hybrid Forecasting with FA

And just as in Section 5.2.1, we repeat these experiments with a different estimate of passenger WTP for Airline 1. As we previously did, we replace the moderate FRAT5 series "C" with the more aggressive "A4" in order to gauge the performance of both airlines when AL1 uses HF and universal FA together, but with a very aggressive estimate of WTP.

For this scenario, we must present a new "No FA Base" – one in which Airline 1 uses DAVN with HF only, and estimates WTP using FRAT5 series "A4". This scenario was previously presented in Section 4.3, and is summarized in Table 35.

Table 35: No FA Base Case Results – DAVN w/ HF versus DAVN, FRAT5 Series "A4"

Airline	Revenue	Load Factor (%)	Yield ($/RPM)
AL1	$1,076,200	85.09	0.1031
AL2	$1,014,869	82.89	0.0961

As previously mentioned, Airline 1's use of FRAT5 "A4" for passenger WTP increases its network revenue to $1,076,200 from $1,072,646 with "C" – a slight increase of 0.33% with the more aggressive series. Again, AL1's use of HF gives it a significant advantage – approximately 6% – over its competitor in terms of network revenue.

When adding Fare Adjustment to this Base Case, we again observe an increase in both airlines' revenue as FA becomes more aggressive, as shown in Figure 61. As in the case of FA alone, the biggest beneficiary to AL1's FA appears to be AL2, whose revenues sharply increase as the FA becomes more aggressive. But unlike with FRAT5 "C", the use of "A4" for both HF and FA shows that there is indeed a limit to the benefit of FA before AL1's revenue starts to decline.

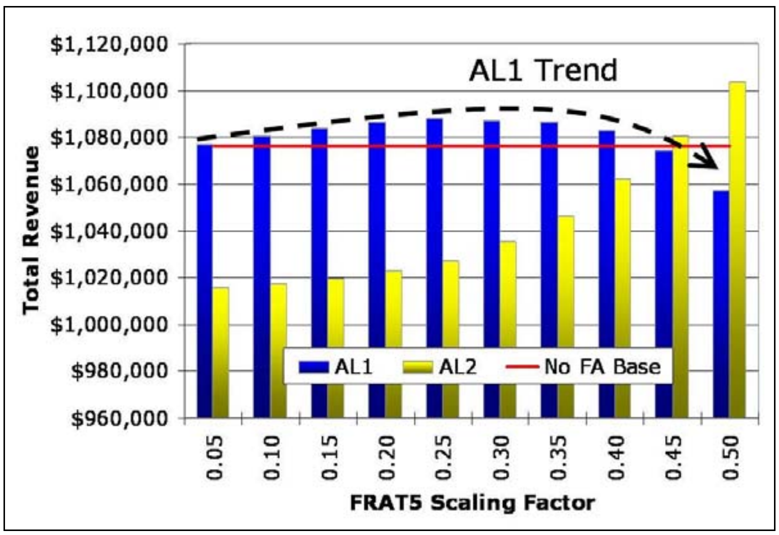

Figure 61: Network Revenue by FRAT5 Scaling Factor for DAVN w/ HF and FA versus DAVN, FRAT5 "A4"

As shown in Table 36, AL1 experiences a peak revenue gain of 1.1% over the HF alone with a FRAT5 Scaling Factor of 0.25. Beyond this point, revenue quickly declines and even leads to a loss from HF alone at Scaling Factors of 0.45 and above.

Table 36: DAVN w/ HF and FA versus DAVN, Airline 1 Change in Network Revenue from DAVN w/ HF by FRAT5 Scaling Factor, FRAT5 "A4"

FRAT5 Scaling Factor	0.05	0.10	0.15	0.20	0.25
Δ AL1 Revenue from HF Alone	0.07%	0.41%	0.71%	0.96%	1.10%
FRAT5 Scaling Factor	**0.30**	**0.35**	**0.40**	**0.45**	**0.50**
Δ AL1 Revenue from HF Alone	1.03%	0.97%	0.64%	-0.18%	-1.77%

Once again, the load factors and yields for both airlines offer clues as to why AL1's revenue peaks when using HF and FA together, while its competitor sees rapid escalation of its own network revenue. As shown in Figure 62, Airline 1's load factor sharply decreases and its yield sharply increases as the FA becomes more aggressive.

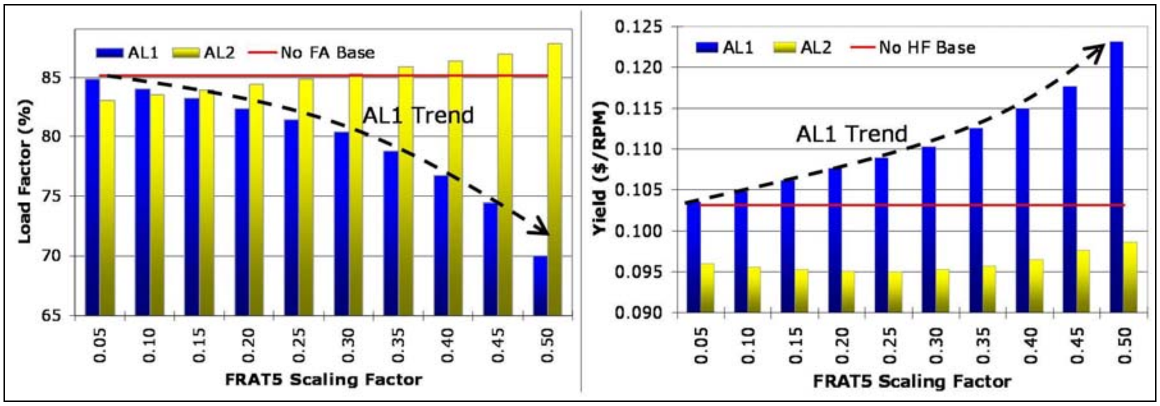

Figure 62: Load Factor and Yield by FRAT5 Scaling Factor for DAVN w/ HF and FA versus DAVN, FRAT5 "A4"

Due to its aggressive WTP estimate of FRAT5 series "A4", when AL1's Scaling Factor grows, the added benefits of FA to HF are outweighed by the growing severity of FA's overprotection problem, as originally described in Section 5.2.1. By the time the Scaling Factor reaches 0.50, the benefits of simultaneous HF disappear, and the overprotection causes load factor to tumble below 70% from a high of 85.09% with HF alone. When AL1's FA becomes this aggressive (not only because of the high scaling factors but because of the aggressive FRAT5 series used), the RM system performs as if no HF at all is being used, and we observe results as such.

Of course, the beneficiary of AL1's overprotection issues is Airline 2, which again observes a gradual elevation of its load factor. But more interesting is that its yield – after slowly declining as with FRAT5 "C" – actually begins to climb upward again when AL1's Scaling Factor reaches 0.30 or thereabouts – precisely the point where AL1's revenues begin to decline. This observation further supports the hypothesis that FA can become too aggressive to help AL1. At this critical point, AL1's FA is so aggressive and denies so many bookings that AL2 can become more selective with its own bookings, thus simultaneously increasing both its yield and load factor, and driving its revenue considerably upward.

5.2.2.3 Summary of Adding HF to FA

So an airline using DAVN in semi-restricted fare structures can not only improve its network revenue by switching from standard pick-up forecasting to HF (see Chapter 4), but can see an even greater benefit when moving from pick-up forecasting with Fare Adjustment to HF with FA, as shown in the above discussion. We have shown the potential increase in AL1's network revenue can be substantial. Dependent upon the Scaling Factor used, the benefit of HF with FA over FA alone approaches 5% for an airline assuming FRAT5 "C" for passenger WTP, and is upwards of 10% when using FRAT5 "A4", as shown in Figure 63. For example, HF with FA (FRAT5 "C") leads to a gain of 3.31% over FA alone at a Scaling Factor of 0.15; at a higher Scaling Factor of 0.25, this gain increases to 3.64%, and continues to grow with the Scaling Factor. Of course, these results are even more pronounced with a more aggressive FRAT5 series.

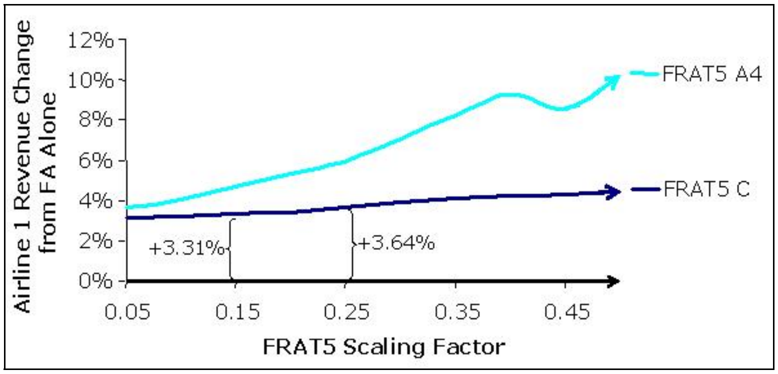

Figure 63: Change in AL1 Network Revenue by FRAT5 Scaling Factor, from DAVN w/ FA to DAVN w/ HF and FA versus DAVN

However, these benefits are extremely misleading when taken out of context. As shown in our experiments, these large revenue gains for adding HF to FA are not necessarily due to the strength of HF, but to the poor performance of FA with pick-up forecasting in semi-restricted fare structures.

In the case of FRAT5 "C", Fare Adjustment without HF leads to a slow deterioration of network revenues as the FRAT5 Scaling Factor becomes more aggressive. When AL1 also uses HF in conjunction, an increase in revenue is observed over the "No FA Base", as discussed in Section 5.2.1.1 (FA, no HF), Section 5.2.2.1 (HF, FA together), and illustrated in Figure 64. As shown in this Figure below, the trends in network revenue diverge as FA becomes more aggressive. So FA with HF appears to continuously help Airline 1 while FA alone continuously hurts it, thus explaining the ever-increasing benefit of HF as FA becomes more aggressive.

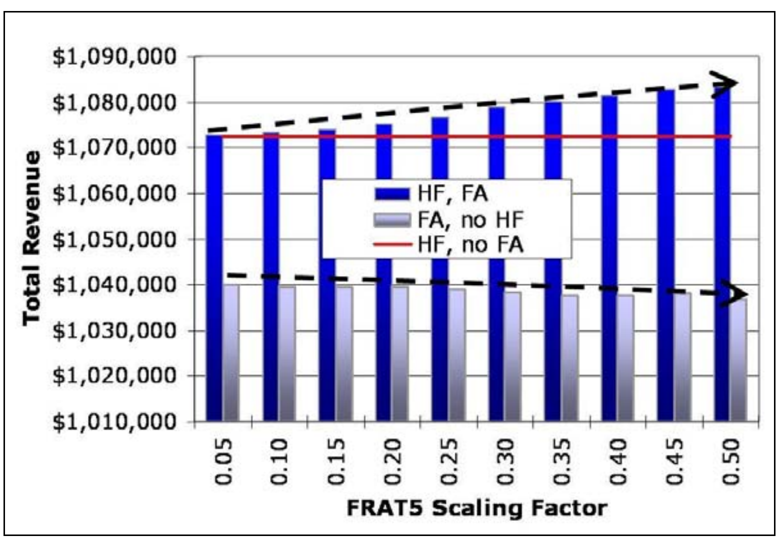

Figure 64: AL1 Network Revenue by FRAT5 Scaling Factor, Comparison of DAVN w/ HF and FA, w/ FA, and w/ HF versus DAVN, FRAT5 "C"

However, this improvement for HF and FA together does not continue indefinitely as the Scaling Factor swells. There exists that critical point where revenues will begin to decline, as demonstrated by switching from input FRAT5 series "C" to the more aggressive "A4". Again, we observe an increase in revenue from FA alone (discussed in Section 5.2.1.2) to the case of HF and FA together (discussed in Section 5.2.2.2), as illustrated in Figure 65. However, AL1's revenues for HF and FA together peak at a Scaling Factor of 0.25, and drop significantly thereafter.

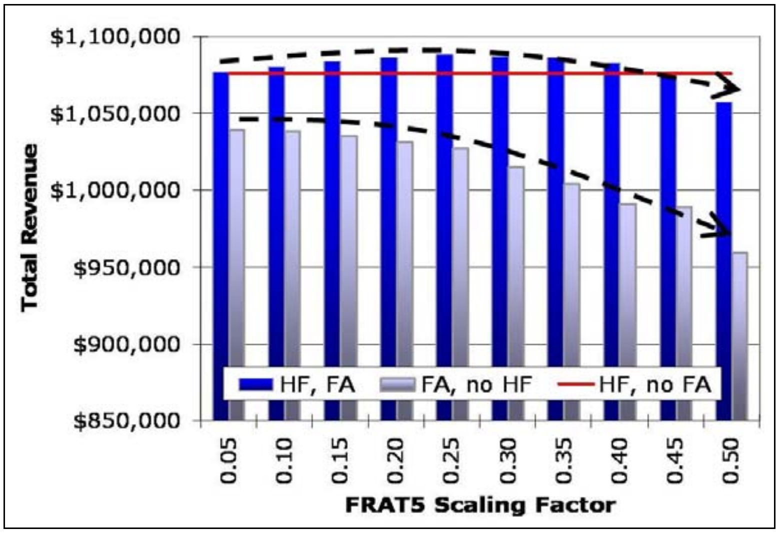

Figure 65: AL1 Network Revenue by FRAT5 Scaling Factor, Comparison of DAVN w/ HF and FA, w/ FA, and w/ HF versus DAVN, FRAT5 "A4"

So as FA becomes too aggressive, both for the case of FA alone and FA and HF together, revenues decline considerably. In the case of FA alone, this deterioration begins immediately, whereas when HF is added to FA, the airline experiences a period of revenue increase before the dramatic decline begins. These large drops in network revenue are the result of overprotection due to FA, and grow in severity as FA becomes more aggressive, as previously discussed. Adding HF to FA leads to gains in total revenue, as long as the FA is not too aggressive. But as shown in our experiments, over-aggressive FA can erase these gains as the overprotection dominates HF's benefits.

While the switch to HF appears to be a potential revenue boon for an airline using FA and pick-up forecasting in semi-restricted fare structures, as shown in Figure 63, that airline should take pause. Our simulations indicate that those revenue gains are largely the result of reversing the losses brought about by stand-alone FA itself.

5.2.2.4 Adding FA to HF

We now examine the complementary situation to Section 5.2.2.3 above. Instead of an airline adding HF to FA, we turn our attention to an airline considering adding FA to HF. Doing so addresses one of the goals this thesis – answering the question "Can Fare Adjustment be used to improve the performance of Hybrid Forecasting in simplified fare structures?" Based on our experiments in this section, the answer is yes.

In addition to FRAT5 series "C" and "A4", we repeated the experiments in Sections 5.2.2.1 and 5.2.2.2, respectively, for the scenarios in which Airline 1 used "A" and "A2" as estimates of passenger WTP. Relative to the appropriate Base Case from Section 4.3 – DAVN w/ HF versus DAVN, no Fare Adjustment – we observe a potential revenue gain for each input FRAT5 series, as shown Figure 66.

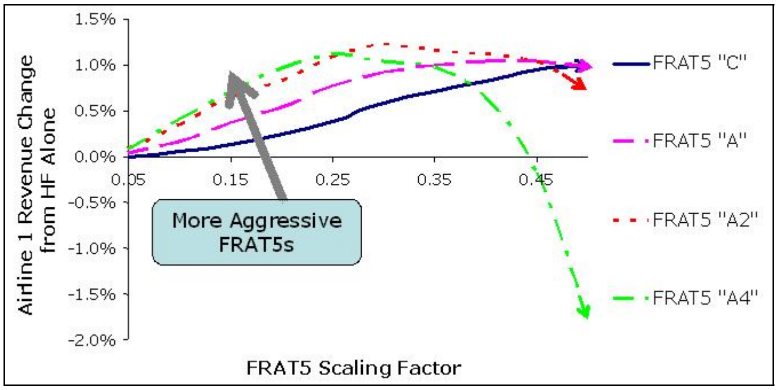

Figure 66: Change in AL1 Network Revenue by FRAT5 Scaling Factor, from DAVN w/ HF to DAVN w/ HF and FA versus DAVN

As the input FRAT5 series becomes more aggressive, the point at which Airline 1's revenue peaks occurs earlier (in terms of the FRAT5 Scaling Factor). For example, the most aggressive WTP estimate (series "A4") has an observed maximum gain of 1.1% at a Scaling Factor of 0.25, as discussed previously. For less aggressive "A2", its peak of 1.21% occurs at a Scaling Factor of 0.30. Similarly, the peak for "A" of 1.04% occurs at 0.40, and the peak for "C" has not been observed due to limiting Scaling Factors to a maximum of 0.50. The revenue increases over the respective Base Cases are shown in Table 37.

Table 37: DAVN w/ HF and FA versus DAVN, Change in AL1 Network Revenue by FRAT5 Scaling Factor from no FA Base Case

FRAT5 Scaling Factor	**0.05**	**0.10**	**0.15**	**0.20**	**0.25**
FRAT5 Series "C"	0.00%	0.05%	0.13%	0.25%	0.39%
FRAT5 Series "A"	0.04%	0.16%	0.36%	0.54%	0.77%
FRAT5 Series "A2"	0.08%	0.35%	0.65%	0.82%	1.08%
FRAT5 Series "A4"	0.07%	0.41%	0.71%	0.96%	1.10%
FRAT5 Scaling Factor	**0.30**	**0.35**	**0.40**	**0.45**	**0.50**
FRAT5 Series "C"	0.59%	0.71%	0.82%	0.95%	1.00%
FRAT5 Series "A"	0.92%	0.99%	1.04%	1.03%	0.96%
FRAT5 Series "A2"	1.21%	1.15%	1.11%	1.03%	0.74%
FRAT5 Series "A4"	1.03%	0.97%	0.64%	-0.18%	-1.77%

This phenomenon occurs because the size of the actual adjustment applied to each fare is dependent both upon the airline's estimate of WTP (FRAT5 series) and its choice of FRAT5 Scaling Factor. So as the aggressiveness of one of these variables increases, the aggressiveness of the other decreases at the point of maximum revenue impact. In our case, as the aggressiveness of Airline 1's input FRAT5 series increases from "C" to "A4", the Scaling Factor for which its network revenue is highest decreases from 0.50 (or above) to an observed value of 0.25.

5.2.2.5 Summary of Fare Adjustment and HF

In conclusion, universal FA - using the same input FRAT5 on all paths regardless of path quality – appears to indeed be a valuable tool for improving the performance of HF in semi-restricted fare structures. By itself, FA tends to hurt network revenues by overprotecting the higher fare classes (as shown in Section 5.2.2.1 and 5.2.2.2). The same can happen in conjunction with Hybrid Forecasting if the FA is applied too aggressively (as shown in Section 5.2.2.3 and 5.2.2.4). However, an airline using moderate FA to supplement its HF can potentially see gains in its network revenues of over 1%, as demonstrated by the simulations in this section.

5.3 Hybrid Forecasting using Path Categories and Fare Adjustment together

In Section 5.1 we demonstrated how the use of Path Categorization to selectively adjust the aggressiveness of Hybrid Forecasting in terms of path quality can improve network revenue by nearly 0.25% over HF alone. And in Section 5.2 (specifically in Section 5.2.2) we showed how the use of moderate Fare Adjustment with Hybrid Forecasting – independent of any Path Categorization – can improve network revenues by over 1% from HF without FA. In this section, we combine PCAT and FA to demonstrate the combined effect both techniques can have for an airline practicing HF with a DAVN optimizer in semi-restricted fare structures.

Similar to the organization of Section 5.2, we first present a series of experiments without Hybrid Forecasting in Section 5.3.1 to demonstrate the performance of combined PCAT and FA without the aid of HF. Then in Section 5.3.2, we repeat these experiments with Airline 1 replacing its standard pick-up forecasting with Hybrid Forecasting. Doing so allows us to best gauge the impact of HF and gain better insight into the performance of combined FA and PCAT.

Because we simultaneously test three techniques in this section – PCAT, FA, and HF – we have multiple dimensions by which to analyze the performance of AL1's Revenue Management system. For this reason, we cannot refer to a single Base Case for comparison in each experiment. Thus, the appropriate Base Case for measuring variation in network revenue will be presented for each simulation where appropriate.

Furthermore, there are numerous variations of PCAT, FA, and HF thanks to the array of WTP estimates (FRAT5 series) and Scaling Factors available to AL1. In an effort to both simplify analysis as well as further isolate the effects of these three techniques, we limit the scope of our experiments in this section in two ways:

1. We restrict Airline 1 to the baseline FRAT5 series "C" – a moderate estimate of passenger WTP given the 17 series used in PODS (as discussed in Section 3.2.2.1 and Section 4.1.3).

2. We restrict the scope of Airline 1's Fare Adjustment FRAT5 Scaling Factors of 0.15, 0.25, or 0.50. As shown in Section 5.2, there is enough spread among these three values to demonstrate trends in FA despite the reduction in our simulation efforts.

5.3.1 DAVN w/ FA and PCAT versus DAVN

As described above, our initial experiments combining Fare Adjustment with Path Categorization exclude Hybrid Forecasting. The simulations presented in Section 5.1 indicate that the use of higher passenger WTP estimates in PCAT1 OD markets leads to revenue gains for an airline using Hybrid Forecasting. Despite our inability to conclude that passenger WTP in PCAT1 indeed exceeds that in PCAT2 and PCAT3, we have an intuitive expectation of such passenger behavior, and the improvement in total revenue, as well as revenue due to passenger sell-up (specifically in the higher fare classes) supports the idea that non-stop service is preferred to connecting.

Whereas we originally experimented with higher WTP in the context of HF, we now apply Path Categorization exclusively to Fare Adjustment, thus departing from the universal FA described in Section 5.2. By using more aggressive FA in PCAT1 OD markets than in PCAT2 and PCAT3, we expect to improve Airline 1's revenues. Within PODS, we can make FA more aggressive for certain OD markets by selectively increasing the WTP estimate (or input FRAT5 series) for those markets – not by increasing the Scaling Factor, which is a universal parameter for all markets served by a simulated airline.

There are two appropriate Base Cases for this situation. The first Base Case is the universal BC1 (originally shown in Chapter 4 and summarized in Table 9) which describes the scenario with neither Fare Adjustment nor Path Categorization. Comparison with BC1 is needed to measure the improvement in Airline 1's revenue using FA and PCAT relative to using neither.

The second (generalized) BC is that of FA without PCAT, as originally presented in Section 5.2.1, and referred to as the "No PCAT Base". Due to the tendency of FA by itself to lose revenue relative to BC1, this second BC is used to measure the improvement in Airline 1's revenue due exclusively to Path Categorization. Of course, the degree of FA (in terms of the Scaling Factor) must be normalized appropriately among all experiments to ensure an appropriate comparison. For example, the "No PCAT Base" Case is summarized in Table 38 for our chosen FRAT5 Scaling Factors of 0.15, 0.25, and 0.50.

Table 38: No PCAT Base Case Results – DAVN w/ FA versus DAVN, FRAT5 Series "C", Scaling Factor of 0.15, 0.25, 0.50

FRAT5 Scaling Factor	Airline	Revenue	Load Factor (%)	Yield ($/RPM)
0.15	AL1	$1,039,592	81.41	0.1041
	AL2	$1,033,529	82.81	0.0980
0.25	AL1	$1,038,967	79.91	0.1060
	AL2	$1,038,091	83.28	0.0978
0.50	AL1	$1,036,786	76.25	0.1108
	AL2	$1,052,509	84.85	0.0974

With a Scaling Factor of 0.15, Airline 1 has lost only 0.03% of its revenue from BC1 without Fare Adjustment. At this stage of mild FA, AL1's revenue still exceeds AL2's by $6,063. When the Scaling Factor increases to 0.25, the gap between the two airlines has closed to only $876; AL1 has lost 0.13% of its BC1 revenue, while its competitor has experienced a revenue gain of 0.73% over its own BC1 level. Under the most aggressive FA scenario here (a Scaling Factor of 0.50), Airline 2's revenue has swelled to $1,052,509 – 2.13% greater than in BC1. Conversely, AL1's revenue falls by 0.34% under this scenario.

So as Airline 1's FA becomes more aggressive, its revenue gradually erodes relative to the case of no FA (BC1) while AL2 experiences significant increases in its own revenue. These trends were previously described in Section 5.2.1.1 for FRAT5 "C", and are attributed to FA's tendency to overprotect in anticipation of future bookings in higher fare classes.

Airline 1's estimate of passenger WTP was then adjusted in the 40 PCAT1 OD markets, as in Section 5.1; the FRAT5 series used on these paths varied from "D" through "A9" for each of the Scaling Factor values described above. As with HF by Path Category, we intuitively expect Airline 1's network revenue to increase as the FRAT5 series used in PCAT1 becomes more aggressive. And as illustrated in Figure 67 for Scaling Factor 0.15, our intuition appears to be valid. AL1 sees elevations in network revenue with more aggressive FA in its dominant markets.

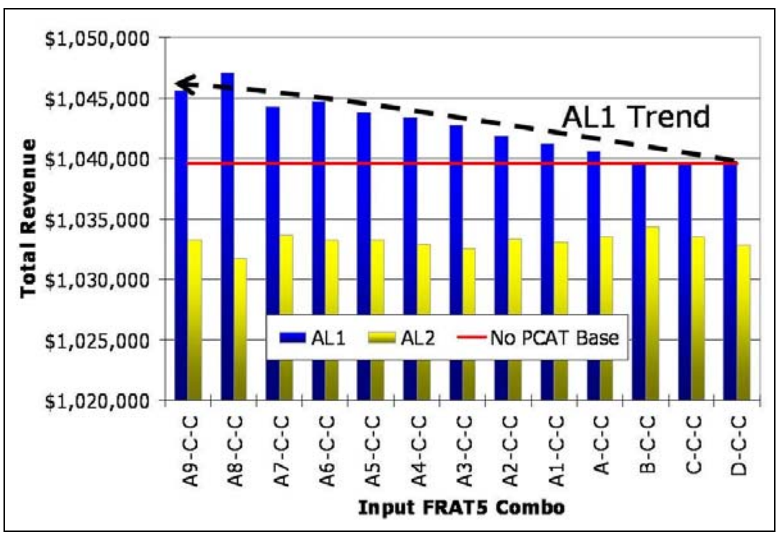

Figure 67: Network Revenue by FRAT5 Series Combination for DAVN w/ FA and PCAT versus DAVN, Scaling Factor of 0.15

And unlike the case of HF by PCAT, there appears to be no general trend in Airline 2's revenue when the aggressiveness of Fare Adjustment is varied by Path Category. These observations regarding each airline's network revenue indicate that FA and PCAT together helps AL1 and does not necessarily help (or hurt) AL2 in a predictable way, in contrast with our experiments with PCAT in Section 5.1.

The revenue changes for Airline 1 relative to the "No PCAT Base" and BC1 are shown in Table 29 for a Scaling Factor of 0.15. Compared to the "No PCAT Base", AL1's network revenue gradually increases to a peak gain of 0.72%, which occurs when FRAT5 "A8" is used on the PCAT1 paths.

Table 39: DAVN w/ FA and PCAT versus DAVN, Airline 1 Revenue Change by PCAT1 FRAT5 Series, Using "C" for PCAT2 and PCAT3

PCAT1 FRAT5 Series	A9	A8	A7	A6	A5	A4	A3
Δ AL1 Revenue from No PCAT Base	0.58%	0.72%	0.45%	0.49%	0.41%	0.37%	0.30%
Δ AL1 Revenue from BC1	0.51%	0.66%	0.38%	0.42%	0.35%	0.30%	0.24%

PCAT1 FRAT5 Series	A2	A1	A	B	C	D
Δ AL1 Revenue from No PCAT Base	0.22%	0.15%	0.09%	-0.01%	-	0.03%
Δ AL1 Revenue from BC1	0.16%	0.09%	0.03%	-0.07%	-0.07%	-0.04%

And with respect to BC1, the use of FA and PCAT together produces revenue gains for all PCAT1 series more aggressive than "B" when FA is used with a Scaling Factor of 0.15. This observation is significant because it indicates that an airline can somewhat correct (or at least offset) the overprotection issue which undermines any revenue improvements due to universal FA alone by augmenting its estimate of passenger WTP in its PCAT1 OD markets. Indeed, it appears that the use of more aggressive FRAT5 series in a relatively small number of markets (only 40 of 482 in PODS Network D-6) can reverse the revenue losses due to FA originally observed in Section 5.2.1.

The change in Airline 1's network revenue relative to the "No PCAT Base" is shown in Figure 68 for Scaling Factors of 0.15 and 0.25. As just discussed, AL1's revenue gradually increases to a peak of 0.72% with a Scaling Factor of 0.15 for input FRAT5 combination "A8-C-C". When more aggressive FA is used by way of increasing the Scaling Factor to 0.25, the same gradual increase in network revenue is again observed. But in this case, the peak revenue gain is only 0.54% above the respective "No PCAT Base"; this peak also occurs earlier with input FRAT5 combination "A2-C-C".

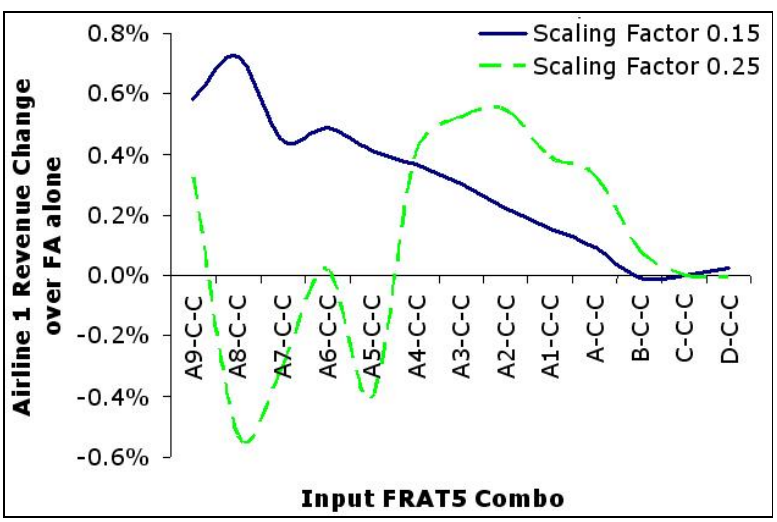

Figure 68: Change in AL1 Network Revenue by FRAT5 Series Combination, from DAVN w/ FA to DAVN w/ FA and PCAT versus DAVN

Also of note is that for WTP estimates beyond FRAT5 series "A4" in PCAT1 (for Scaling Factor 0.25), Airline 1's revenue follows no apparent trend, increasing and decreasing unpredictably as the PCAT FRAT5 series becomes more aggressive. When repeated with a Scaling Factor of 0.50, no trend is observed at all, and the revenue swings from one FRAT5 combination to the next are even more severe than with a Scaling Factor of 0.25; for this reason, the set of FA and PCAT experiments using a Scaling Factor of 0.50 is not shown in Figure 68.

These observations indicate – as in Section 5.2 – that the gradual revenue gains due to more aggressive FA eventually peak and then quickly deteriorate as the FA becomes too aggressive. For the case of PCAT and FA together, our simulations indicate that beyond a certain point (for example "A4-C-C" with a Scaling Factor of 0.25) AL1's FA becomes too aggressive and produces unpredictable changes in network revenue. Beyond this critical point, the randomness in PODS which describes passenger arrival and attempted booking behavior (see Section 3.1) becomes amplified in terms of network revenue because the FA is so aggressive (due to the input FRAT5 series) that very few bookings can occur.

Also encouraging for the combination of PCAT and FA is the observed change in Airline 1's network revenue relative to BC1, or the absence of both PCAT and FA. These changes for Scaling Factors of 0.15 and 0.25 are shown in Figure 69, which is very similar to Figure 68 (the relative increase over FA alone). This similarity is due to the relatively small revenue losses between FA alone and BC1 with this combination of FRAT5 series and Scaling Factor ("C" and 0.15, 0.25), as discussed previously.

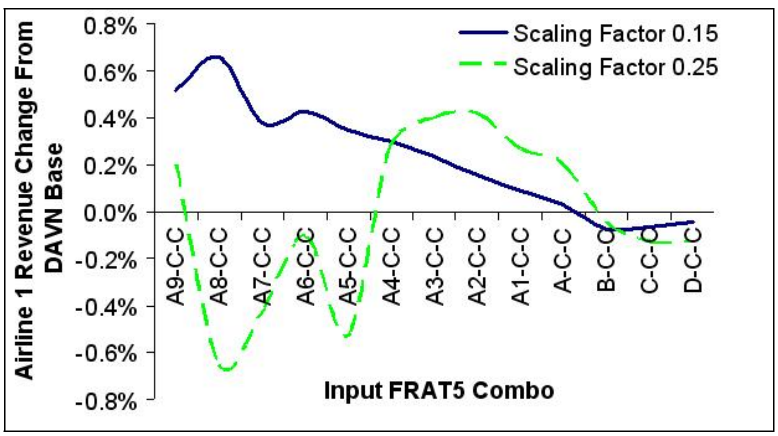

Figure 69: Change in AL1 Network Revenue by FRAT5 Series Combination, from DAVN (Base Case 1) to DAVN w/ FA and PCAT versus DAVN

Of note here is that PCAT and FA together can lead to revenue gains for Airline 1 beyond 0.5% over BC1. Because of the losses of stand-alone FA, our observed revenue improvement over BC1 via FA and PCAT together further emphasizes the value of Path Categorization by relative path quality – a point discussed in Section 5.1.

In summary, our experiments with Fare Adjustment in conjunction with Path Categorization support the following three statements:

1. When added to FA alone, PCAT can improve network revenues by capturing more sell-up from lower to higher fare classes – an expected result given passenger preference for non-stop service.

2. In general, the network revenue for an airline using FA and PCAT together grows as the FA becomes more aggressive (either by way of more aggressive PCAT1 FRAT5 series or higher Scaling Factors). But past a critical point, the FA becomes too aggressive, restricts too many bookings, and leads to unpredictable and/or undesirable changes in network revenue.
3. While FA alone tends to erode revenues relative to no FA (see Section 5.2.1), FA and PCAT together can increase an airline's revenue by approximately 0.5%. This supports our previous belief (see Section 5.1) that the value in Path Categorization by path quality lies in the significant revenue improvements made possible by increasing WTP estimates for a relatively small number of paths.

5.3.2 DAVN with HF, FA, and PCAT together versus DAVN

Having demonstrated the potential revenue improvements for an airline using simultaneous Fare Adjustment and Path Categorization, we now demonstrate the performance of an airline RM system using HF, FA, and PCAT in conjunction. To do so, we repeat the experiments presented in Section 5.3.1 – FA and PCAT using FRAT5 series "C" with Scaling Factors of 0.15, 0.25, and 0.50, and series "D" through "A9" for PCAT1 – but with Airline 1 now using Hybrid Forecasting instead of pick-up forecasting.

As in that previous section, one of the appropriate Base Cases for comparison – of which there are now several – is the "No PCAT Base", but with AL1 also using HF. This scenario was previously presented in Section 5.2.2, and is summarized in Table 40 for our three FRAT5 Scaling Factors.

Table 40: No PCAT Base Case Results – DAVN w/ HF and FA versus DAVN, FRAT5 Series "C", Scaling Factor of 0.15, 0.25, 0.50

FRAT5 Scaling Factor	Airline	Revenue	Load Factor (%)	Yield ($/RPM)
0.15	AL1	$1,074,040	85.83	0.1020
	AL2	$1,015,585	81.69	0.0976
0.25	AL1	$1,076,800	85.25	0.1030
	AL2	$1,017,129	82.15	0.0972
0.50	AL1	$1,083,398	83.35	0.1060
	AL2	$1,023,682	83.36	0.0964

As demonstrated in Section 5.2.2.1 and presented in Table 40, Airline 1's use of the moderate FRAT5 series "C" as an estimate of passenger WTP leads to increases in network revenue for both carriers as Fare Adjustment becomes more aggressive.

Relative to this "No PCAT Base", we observe an increase in Airline 1's revenue as its PCAT1 FRAT5 series becomes more aggressive, as shown in Figure 70 for a Scaling Factor of 0.15. Due to the efficacy of Path Categorization to improve revenue as previously demonstrated in Section 5.1 and Section 5.3.1, this gain is expected for Airline 1. But the revenue gains over the "No PCAT Base" are far greater than those seen over either HF alone or FA alone. With FRAT5 series combination "A9-C-C" – the most aggressive tested – AL1 improves its revenues by 1.52% over that of the "No PCAT Base" of "C-C-C".

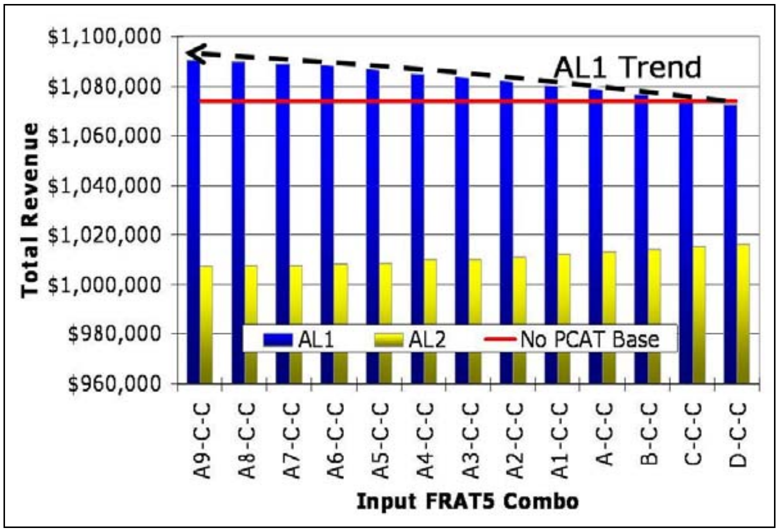

Figure 70: Network Revenue by FRAT5 Series Combination for DAVN w/ HF, FA, and PCAT versus DAVN, Scaling Factor of 0.15

As shown in Figure 71, this upward trend over the "No PCAT Base" continues for both Scaling Factors of 0.25 and 0.50. In the case of the Scaling Factor 0.25, FRAT5 combination "A9-C-C" again leads to the highest observed revenue gain over the "No PCAT Base" – now 1.86%. Note that for the two smallest Scaling Factors (0.15 and 0.25) even larger gains in revenue than those observed are likely because no peak had yet been observed at the most aggressive FRAT5 series combination tested ("A9-C-C").

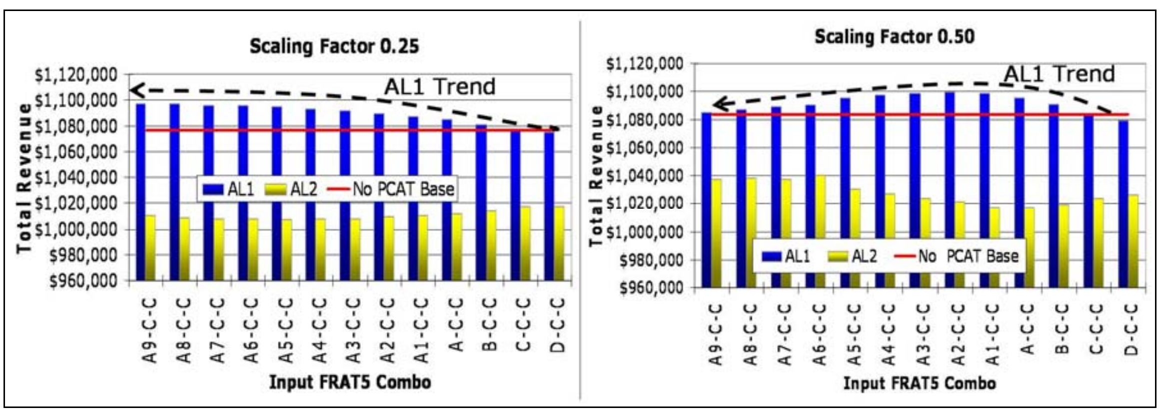

Figure 71: Network Revenue by FRAT5 Series Combination for DAVN w/ HF, FA, and PCAT versus DAVN, Scaling Factor of 0.25 and 0.50

This was not the case, however, for the largest, most aggressive Scaling Factor tested of 0.50. Here, AL1's network revenue peaked at 1.51% over the "No PCAT Base" with a combination of "A2-C-C". The observation of such a peak again indicates that there is a limit to the benefit of PCAT. When combined with HF and FA, that limit is reached when the combination of forecasting and Fare Adjustment becomes so aggressive in determining booking limits and closing fare classes that AL1 can no longer extract further revenue from its network.

Regarding Airline 1's competitor, for each of these three Scaling Factors we observe that AL2's revenue slightly decreases as AL1's improves – an intuitive trend representing the trade-off between two carriers in a highly competitive network. And in the case of Scaling Factor 0.50, AL2's revenue starts to increase once AL1 passes the critical point where its revenues peak. To further understand each airline's changes in revenue with different combinations of AL1's input FRAT5 series, we examine load factors and yields.

More specifically, we inspect each carrier's LF and yield at each FRAT5 combination used by Airline 1, as shown in Figure 72 for a Scaling Factor of 0.15. Remarkably, there is very little change in load factor – and certainly no trend – for either AL1 or AL2 as AL1's PCAT becomes more or less aggressive. Despite the complicated interaction of Hybrid Forecasting and Fare Adjustment in the different Path Categories, Airline 1's RM system manages to keep approximately the same number of bookings over a wide range of WTP estimates in the 40 PCAT1 OD markets. And because its LF is relatively constant in the face of increasing revenue, AL1's yields also increase with more aggressive PCAT (combined with FA and HF, of course). These observations indicate successful capturing of sell-up by Airline 1.

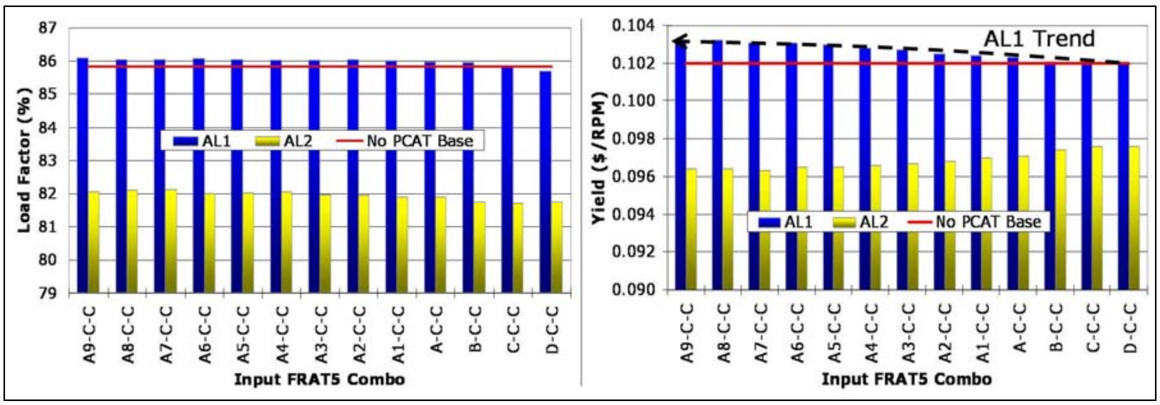

Figure 72: Load Factor and Yield by FRAT5 Series Combination for DAVN w/ HF, FA, and PCAT versus DAVN, Scaling Factor of 0.15

In summary, we have briefly demonstrated the benefits of using HF, FA, and PCAT together by way of presenting a simple example. Though the aggressiveness of the Path Categorization was varied for this example, we just as easily could have examined its sensitivity to input FRAT5 series (in the same manner as Section 4.3) or tested the sensitivity to a varying FRAT5 Scaling Factor (as in Section 5.2.2). With so many variables now in play, there are numerous tests and sensitivity analyses possible.

We focus on the following two scenarios in an effort to further understand the impact of Hybrid Forecasting on an airline RM system, as well as the incremental impact of Fare Adjustment and Path Categorization on Hybrid Forecasting:

1. In Section 5.3.2.1, we examine the incremental revenue gain when adding Hybrid Forecasting to simultaneous Fare Adjustment and Path Categorization.
2. In Section 5.3.2.2, we examine the incremental revenue gain when adding Fare Adjustment and Path Categorization together to Hybrid Forecasting.

5.3.2.1 Adding HF to FA and PCAT

By comparing our simulations of simultaneous Hybrid Forecasting, Fare Adjustment, and Path Categorization with those of FA and PCAT alone (see Section 5.3.1), we observe a potential revenue increase approaching 5% due to the addition of HF, as shown in Figure 73 for Scaling Factors of 0.15 and 0.20. And as the PCAT becomes more aggressive, the revenue gain over FA and PCAT alone tends to increase.

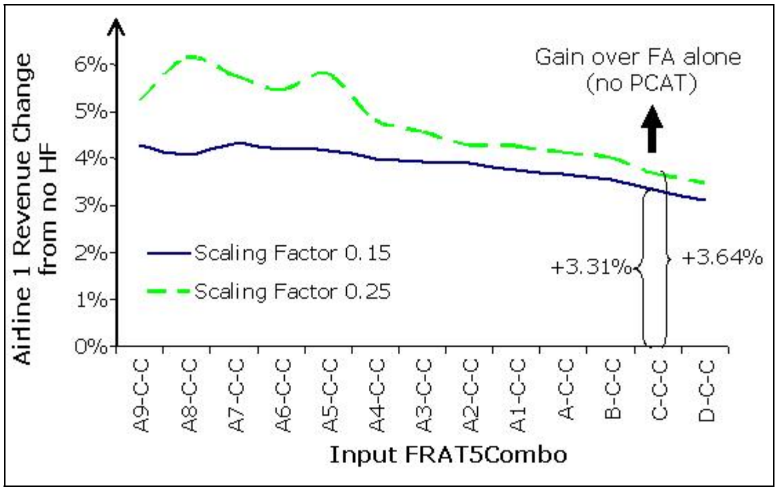

Figure 73: Change in AL1 Network Revenue by FRAT5 Series Combination, from DAVN w/ FA and PCAT to DAVN w/ HF, FA, and PCAT versus DAVN

At a Scaling Factor of 0.15, we observe a revenue gain of 3.31% for the "C-C-C" combination; this gain grows to 3.64% at a Scaling Factor of 0.25. Because the same FRAT5 series is used for all 482 OD markets, these values represent the gain due to HF over FA alone. Note that the symmetry with Figure 63 in Section 5.2.2.3 – our discussion of the improvements HF can produce over FA alone.

In Figure 74 for Scaling Factor 0.15, we show the overall revenue trend from HF, FA, and PCAT together (seen before in Figure 70) as well as the trend from FA and PCAT together (seen before in Figure 67) with respect to the FRAT5 series combination. The difference between these two trends represents the revenue gain due to HF. Because the HF, FA, and PCAT trend increases faster than that of the FA and PCAT, we observe the upward trend for Scaling Factor 0.15 in Figure 73. In the same manner, the revenue benefit of adding HF also grows with more aggressive PCAT for Scaling Factor 0.25.

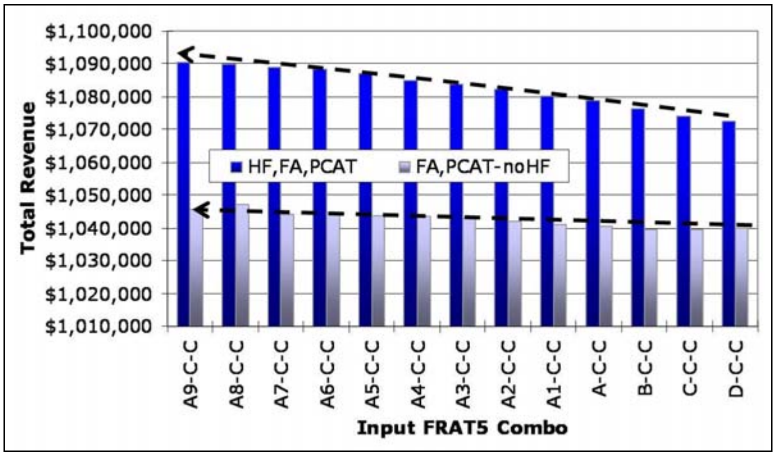

Figure 74: AL1 Revenue by FRAT5 Series Combo., Comparison of DAVN w/ HF, FA, & PCAT and DAVN w/ FA & PCAT versus DAVN, Scaling Factor 0.15

Also, note that Figure 73 omits Scaling Factor 0.50, though not because no improvement due to HF is observed with such an aggressive Scaling Factor. To the contrary, adding HF to FA and PCAT at this larger Scaling Factor still leads to a significant revenue gain at all FRAT5 series combinations, as shown in Figure 75. However, the magnitude of that gain follows no trend due to the lack of a revenue trend for FA and PCAT alone, as previously discussed in Section 5.3.1.

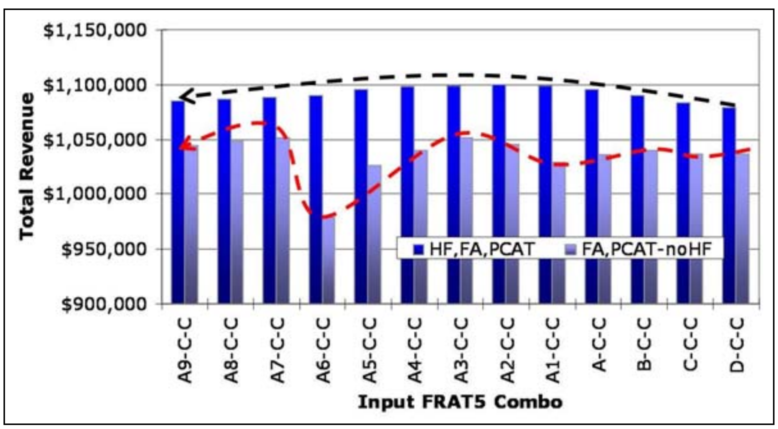

Figure 75: AL1 Revenue by FRAT5 Series Combo., Comparison of DAVN w/ HF, FA, & PCAT and DAVN w/ FA & PCAT versus DAVN, Scaling Factor 0.50

So switching from standard pick-up forecasting to HF appears to significantly increase network revenues for an airline using FA and PCAT together. This is consistent with the jump in revenues previously seen with the switch from pick-up to Hybrid Forecasting in Section 4.3 and Section 5.2.2.3 overall gain, and further strengthens our belief that HF consistently outperforms pick-up forecasting in semi-restricted fare structures.

5.3.2.2 Adding FA and PCAT to HF

Finally, we focus on the incremental benefit that simultaneous PCAT and FA can bring to an airline using HF alone. Here we compare our simulations of HF, FA, and PCAT together with our original HF experiments from Section 4.3. As shown Figure 76, the addition of FA and PCAT appears to increase AL1's revenues over stand-alone HF for all combinations of FRAT5 input series more aggressive than "C" for PCAT1.

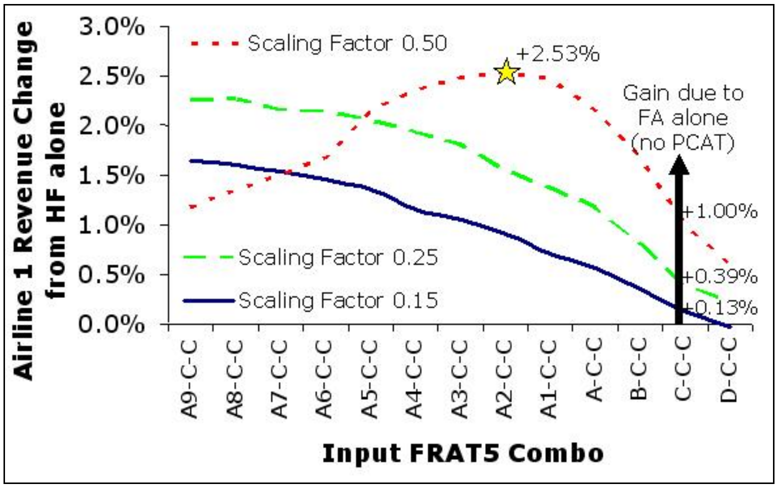

Figure 76: Change in AL1 Network Revenue by FRAT5 Series Combination, from DAVN w/ HF to DAVN w/ HF, FA, and PCAT versus DAVN

Note that at the "C-C-C" combination AL1 sees a small revenue gain over DAVN alone of 0.13% at Scaling Factor 0.15; this value grows to 0.39% and 1.00% at Scaling Factors 0.25 and 0.50, respectively. Because FRAT5 series "C" is applied to each of Airline 1's 482 OD markets (i.e. no Path Categories are being used), these gains represent the revenue increase the revenue gains due solely to FA – consistent with our previous discussion of adding FA to HF alone in Section 5.2.2.4.

As shown in Table 41, the largest observed revenue gain for the mildest Scaling Factor tested of 0.15 – 1.65 % - occurs with our most aggressive FRAT5 series "A9" applied to the PCAT1 markets, indicating that this Scaling Factor's peak gain was not observed.

Table 41: DAVN w/ HF, FA, and PCAT versus DAVN, Change in AL1 Network Revenue by PCAT1 FRAT5 Series from DAVN w/ HF Alone

PCAT1 FRAT5 Series	A9	A8	A7	A6	A5	A4	A3
FRAT5 Scaling Factor 0.15	1.65%	1.60%	1.53%	1.46%	1.36%	1.16%	1.04%
FRAT5 Scaling Factor 0.25	2.25%	2.26%	2.15%	2.14%	2.05%	1.93%	1.81%
FRAT5 Scaling Factor 0.50	1.16%	1.35%	1.50%	1.67%	2.15%	2.35%	2.47%

PCAT1 FRAT5 Series	A2	A1	A	B	C	D
FRAT5 Scaling Factor 0.15	0.91%	0.71%	0.56%	0.35%	0.13%	-0.03%
FRAT5 Scaling Factor 0.25	1.55%	1.37%	1.17%	0.81%	0.39%	0.22%
FRAT5 Scaling Factor 0.50	2.53%	2.47%	2.15%	1.68%	1.00%	0.60%

A peak in revenue improvement was observed for Scaling Factor 0.25 at combination "A8-C-C" (2.26%), and also for Scaling Factor 0.50 at combination "A2-C-C" (2.53%).

Based on these results, it appears that we can indeed improve the performance of HF by simultaneously applying FA and PCAT. As shown in our simulations, Airline 1's RM system saw increases in revenue approaching 2% to 2.5% by supplementing Hybrid Forecasting with Fare Adjustment, and using more aggressive WTP estimates in its dominant OD markets.

5.4 Chapter Summary: Using Path Categories and Fare Adjustment with Hybrid Forecasting

In this chapter we have presented the results of simulations focusing on Path Categorization by relative path quality and Fare Adjustment both with and without Hybrid Forecasting. And based on the results of our simulations, it appears that the performance of Hybrid Forecasting in terms of network revenue can significantly be improved through application of PCAT and/or FA.

In Section 5.1, we demonstrated that passenger WTP seems to be higher for non-stop service compared to connecting options, and we showed how an airline can exploit this elevated WTP in its dominant markets to improve its revenue. Our experiments indicated that the gain over HF alone can approach 0.25%.

Next in Section 5.2, we applied the technique known as Fare Adjustment within an airline's RM system (specifically within the seat allocation optimizer) to test its potential for capturing sell-up. While FA alone proved unsuccessful, and actually lost money from our Base Case, FA combined with HF improved revenues. We demonstrated that not only can adding HF to FA improve FA's performance (even to the point of revenue gains), but we also showed that supplemental FA can improve an airline's revenues by approximately 1% over HF alone.

In Section 5.3, we combined the PCAT and FA techniques of this chapter's two previous sections. Our experiments here showed that using more aggressive FA on an airline's dominant paths can actually reverse the revenue losing performance of regular FA (i.e. applied in the same manner to all OD markets). Furthermore, we showed that replacing standard pick-up forecasting with HF can significantly improve the performance of FA and PCAT together. Finally, we demonstrated that supplemental application of PCAT and FA can improve an airline's revenues by up to 2.5% over HF alone – a significant increase.

To summarize the revenue impacts of HF, FA, and PCAT, we select a "Representative Case" from the set of experiments presented in this chapter (as well as Chapter 4). In this example, Airline 1 uses an input WTP estimate of FRAT5 "C" for its universal HF and/or FA. When practicing a form of Path Categorization (HF, FA, or both), AL1 uses the more aggressive "A2" in its PCAT1 markets. And when practicing FA, Airline 1 uses a Scaling Factor of 0.50.

As shown in Table 42, Airline 1's network revenue improves by 3.11% when switching from pick-up forecasting to Hybrid Forecasting. When practicing FA, AL1 sees a gain of 4.50% due to the switch from pick-up to HF, and when also using "A2-C-C" PCAT, the revenue improvement due to the HF switch grows to 5.22%.

Table 42: Airline 1 Revenue Improvement Due to HF, Representative Case

AL1 RM System	Reference Section	AL 1 Revenue	Δ Revenue With HF	
DAVN w/ pick-up (BC1)	4.1.3	$1,040,277		
…w/ HF	4.3	$1,072,646	$32,369	3.11%
DAVN w/ pick-up and FA	5.2.1	$1,036,786		
…w/ HF	5.2.2	$1,083,398	$46,612	4.50%
DAVN w/ pick-up, FA, PCAT	5.3.1	$1,045,249		
…w/ HF	5.3.2	$1,099,791	$54,542	5.22%

Having demonstrated the potential revenue improvement by switching from pick-up to Hybrid Forecasting in semi-restricted fare structures, we now present the potential revenue improvement over HF alone due to FA and PCAT. Focusing on our representative case, Table 43 shows the improvement in Airline 1's network revenues for three scenarios. By using HF varied by relative path quality ("A2-C-C"), AL1 improves revenue slightly by 0.16%. When instead using FA with a Scaling Factor of 0.50, the improvement over HF alone is 1%. And when combining FA and PCAT, Airline 1 sees a gain of 2.53%.

Table 43: Airline 1 Revenue Improvement over HF, Representative Case

AL1 RM System	Reference Section	AL1 Revenue	Δ Revenue from HF Alone	
DAVN w/ HF	4.3	$1,072,646	-	-
…w PCAT	5.1	$1,074,350	$1,704	0.16%
…w FA	5.2.2	$1,083,398	$10,752	1.00%
…w FA, PCAT	5.3.2	$1,099,791	$27,145	2.53%

Finally, the overall potential revenue gains over standard pick-up forecasting (Base Case 1) in our representative case are shown in Table 44. As described previously in this chapter, universal FA without the aid of HF leads to a loss of revenue (0.34% in this representative case) but the use of PCAT can help (0.48% over pick-up forecasting here). The use of HF with PCAT improves Airline 1's revenue by 3.28%, HF with FA improves it by 4.15%, and HF, FA and PCAT together improves it by 5.72% over pick-up forecasting.

Table 44: Airline 1 Revenue Improvement over Pick-up Forecasting, Representative Case

AL1 RM System	Reference Section	AL1 Revenue	Δ Revenue from BC1	
DAVN w/ pick-up (BC1)	4.1.3	$1,040,277	-	-
…w/ pick-up and FA	5.2.1	$1,036,786	-$3,491	-0.34%
…w/ pick-up, FA, PCAT	5.3.1	$1,045,249	$4,972	0.48%
…w/ HF	4.3	$1,072,646	$32,369	3.11%
…w/ HF & PCAT	5.1	$1,074,350	$34,073	3.28%
…w/ HF & FA	5.2.2	$1,083,398	$43,121	4.15%
…w/ HF, FA, PCAT	5.3.2	$1,099,791	$59,514	5.72%

6 Conclusions

6.1 Summary of the Problem and Thesis Objectives

We began this thesis by introducing a problem specific to the air transportation industry: the effectiveness of the sophisticated Revenue Management (RM) systems employed by many airlines has weakened, ostensibly due to the growth of Low Cost Carriers (LCC). Traditionally these RM systems have relied on a set of booking restrictions (Saturday night stay requirements, non-refundability, etc.) to fence potential demand into well-segregated fare classes. By limiting passenger eligibility for these fare classes, airlines had segregated their passengers and steered them into booking specific classes – effectively making the passenger demand among various fare classes independent. This independence assumption is at the core of traditional RM, as described in Chapter 1. And it is this independence assumption which is violated when those fare class restrictions are lessened, requiring the development of new RM techniques for use in less-restricted environments.

The primary technique examined in this thesis is Hybrid Forecasting (HF), and is a marriage of forecasting techniques designed to simultaneously forecast demand for two groups:

1. Product-oriented passengers who seek a specific fare class when booking despite the potential availability of other classes (which may be lower priced);
2. Price-oriented passengers who simply aim to book in the lowest available fare class for which they are eligible. It is these passengers who dilute revenues in traditional RM systems in less-restricted environments due to the availability of multiple homogenous fare classes indistinguishable except for price.

As introduced in the opening chapter, the primary goal of this thesis is to measure the potential revenue gains due to HF in semi-restricted fare environments. Beyond this objective, we also seek to measure potential revenue improvements over HF by using two supplemental techniques:

1. Fare Adjustment (FA) – a tool developed to encourage passenger sell-up by having the seat allocation optimizer proactively close lower fare classes;
2. Path Categorization (PCAT) – the selective application of higher passenger willingness-to-pay (WTP) estimates in Origin to Destination (OD) markets where path quality dominates competitors (i.e. non-stop versus connecting);

In Chapter 2 we presented a review of the relevant literature in order to understand the scope of our problem. We began with a general introduction of passenger demand forecasting in the airline industry, followed by a more focused discussion of RM systems and their components, including the widely used method of traditional pick-up forecasting. Next, we reviewed the emergence of LCCs, and how their simplified fare structures can diminish the performance of traditional RM systems by violating that critical demand independence assumption. We continued by introducing Q-forecasting – a technique for estimating demand in totally unrestricted fare structures (as opposed to our semi-restricted ones) – as well as Fare

Adjustment. Finally, we presented the idea of price-oriented and product-oriented demand, and introduced the concept of a "hybrid" method of demand forecasting to separately estimate both of them for the use of a traditional seat optimizer.

To achieve the goals of this thesis, we simulated HF (as well as FA and PCAT) using the Passenger Origin-Destination Simulator (PODS), a tool introduced in Chapter 3. We first presented the background of PODS, including its passenger choice model and its RM system for each simulated airline. Next, we explained the methodology within PODS used by each of our tools for dealing with less-restricted fare structures, including Q-forecasting (the price-oriented component of HF), FA, and Path Categorization. Finally, we presented nine different ways in which Hybrid Forecasting can categorize historical bookings as product-oriented or price-oriented within the simulator.

6.2 Summary of Findings

As previously described, we used PODS to simulate the performance of HF (as well as FA and PCAT) in a two-airline environment where both carriers use the Displacement Adjusted Virtual Nesting (DAVN) seat allocation optimizer. We began Chapter 4 by introducing this simulated network (Network D-6) as well as the semi-restricted fare structure shared by both airlines. In our experiments, Airline 1 (AL1) used various forms of HF while Airline 2 (AL2) always used standard pick-up forecasting.

Our first set of experiments tested our nine classifications of product-oriented and price oriented demand by simulating each under several scenarios, including with low and high demand, with a leg-based seat optimizer, and with a more aggressive estimate of passenger WTP. In each of these scenarios, one booking classification method outperformed the other eight. When using the "HF1-IAP0" combination, we classify product-oriented and price-oriented bookings as follows:

- All bookings in which the next lowest fare class was available on the same path (same flight and airline) are product-oriented.
- All other bookings, including those made when the next lower class has been closed due to AP requirements, are price-oriented.

Of the nine classifications available within PODS, we demonstrated that HF1-IAP0 appears to outperform the others because method HF1 minimizes revenue loss due to overprotection of the highest fare classes (unlike HF2 and HF3) and method IAP0 minimizes revenue loss due to dilution and spiral-down (unlike IAP1 and IAP2). Based on the results of our experiments in Section 4.2, we adopted HF1-IAP0 as the standard historical booking classification method used for Hybrid Forecasting in this thesis.

In Section 4.3, we tested the sensitivity of HF to the input estimate of WTP used by an airline. In PODS, a simulated airline uses a "FRAT5" value (as described in Section 3.2.2.1) as a proxy for WTP which governs the assumed behavior of passengers selling up from low to high fare classes. By varying Airline 1's input FRAT5 series, we showed that HF with low estimates of WTP can lead to large losses in revenue compared to the use of pick-up forecasting, while more aggressive assumptions of passenger WTP can improve network revenues between 3% and 4%.

After demonstrating the potential value of replacing pick-up forecasting with Hybrid Forecasting in semi-restricted fare structures, we focused on improving the performance of HF in Chapter 5. We began in Section 5.1, by applying more aggressive WTP estimates (in terms of FRAT5 values) in OD markets where a given airline demonstrated dominant path quality (a technique referred to as Path Categorization, or PCAT). Doing so resulted in revenue gains approaching 0.25% over HF with "universal" FRAT5 inputs for all paths. While such a change seems small (especially compared to the 3-4% gains of universal HF over pick-up forecasting), the improvement is actually significant due to the relatively small number of dominant paths in which an airline augments its WTP estimate.

We continued in Section 5.2 by experimenting with universal Fare Adjustment, first simulating FA without HF and demonstrating its inability to improve network revenues over pick-up forecasting. Then we showed the significant revenue gains possible when combining universal HF and FA – greater than 1% over HF alone.

And in Section 5.3 we combined FA and PCAT, which led to even larger revenue gains than with either technique separately. We first showed that varying the aggressiveness of FA by relative path quality led to revenue gains over universal FA – sometimes large enough to improve over pick-up forecasting alone. Finally, we combined HF, FA, and PCAT and demonstrated an improvement over universal HF alone of up to 2.5%.

We summarized the possible revenue changes in each of these above scenarios by presenting a "Representative Case" of an airline switching from pick-up forecasting to Hybrid Forecasting. In this example, Airline 1 used FRAT5 "C" for its universal FA and HF – a moderate value thought to be realistic in terms of passenger WTP. To the extent the airline employed Path Categorization it assumed a more aggressive WTP of FRAT5 "A2" in its dominant OD markets. And when performing Fare Adjustment, Airline 1 used a Scaling Factor of 0.50. Each of the above scenarios involving HF, FA, and PCAT is illustrated in Figure 77 in terms of Airline 1's network revenue.

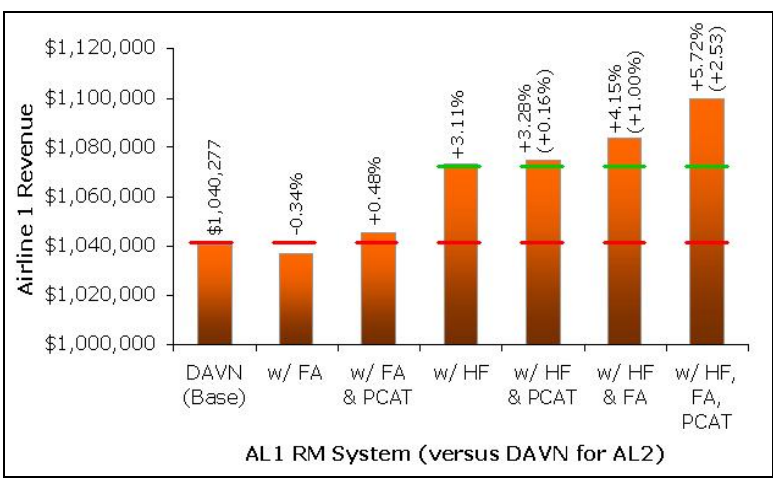

Figure 77: Network Revenue and Changes from Pick-up Forecasting, Airline 1, Representative Case

Relative to the use of pick-up forecasting alone, universal FA decreased revenue by 0.34%, while adjusting the aggressiveness of FA by relative path quality improved revenues by 0.48%. The introduction of HF led to much larger increases in revenue – 3.11% over the use of pick-up forecasting. Furthermore, adding PCAT improved revenues over universal HF by 0.16%, adding FA by 1.00%, and adding both simultaneously by 2.53%. For each of these scenarios in our representative case, the gain over pick-up forecasting was 3.28%, 4.15%, and 5.72%, respectively.

6.3 Future Research Directions

Having demonstrated the potential value of employing Hybrid Forecasting in airline Revenue Management systems with semi-restricted fare structures, we now suggest two related avenues of interest ripe for future investigation. First, it is important to note that the experiments in this thesis – while numerous – were certainly not exhaustive or even conclusive.

As a demonstration of the concept of Hybrid Forecasting in semi-restricted fare structures, we have limited our simulations to two competing airlines with overlapping hub-and-spoke networks, one-way flows, and similar Revenue Management systems. In this specific situation, we have clearly demonstrated the potential revenue benefit afforded to an airline using HF when its lone competitor uses traditional pick-up forecasting. While a valuable initial step in understanding the impact of HF, our two carrier case greatly magnifies the competitive aspects of the airline industry. As demonstrated in many of our experiments, a loss of passengers or revenue by one carrier typically led to a respective gain for the other due to the lack of alternatives as well as the similarity of the connecting service provided in nearly every OD market. To better simulate the performance of HF, we could implement an expanded PODS network with additional competitors in more markets, such as the larger, more competitive "Network R", as described by Dar[78].

Also interesting would be an expansion of the experiments performed here to include cases of competing airlines both using HF as well as with different combinations of seat inventory optimizers. So beyond the "DAVN with HF versus DAVN with pick-up forecasting" case in this thesis, a more thorough evaluation could include the following: "DAVN with HF versus DAVN with HF", or "versus the leg-based EMSRb optimizer" (discussed in Section 3.1.3.3), or versus any number of widely used RM techniques. It is likely that the gains presented demonstrated in this thesis would be diminished if the techniques were used by competing carriers.

Beyond extending the scope of simulations performed to better understand HF, the second suggested research direction is development of a method to estimate passenger WTP. For all experiments in this thesis, we have assumed various levels of passenger WTP – which take the form of FRAT5 values in PODS – in order to manage sell-up behavior. Doing so complicates both experimentation and analysis by necessitating the need for simulation at various assumed WTP levels. Furthermore, the absence of an estimate of passenger WTP hinders the applicability of HF in the airline industry.

---

[78] Dar, M. 2006. Modelling the performance of revenue management systems in different competitive environments. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

In his thesis testing Q-forecasting in totally unrestricted fare structures, Cléaz-Savoyen[71] used a FRAT5 estimation process which performed relatively well. But as described by Vanhaverbeke[76], this estimator is far from perfect. Specifically, this FRAT5 estimator should better vary WTP estimates in response to competitor actions as well as with decreasing days to departure.

But beyond these challenges, our experiments illustrate the need for further estimation beyond that of network-wide FRAT5 values. For use in Fare Adjustment, our simulations were repeated with ten different Scaling Factors applied to the normal FRAT5 values. Clearly, we not only need an estimate of WTP, but the use of Fare Adjustment necessitates some estimate of this Scaling Factor. And unlike with the FRAT5 estimates for passenger WTP, we currently lack even a rudimentary method of estimating this Scaling Factor in PODS.

Furthermore, the revenue improvement in our Path Categorization experiments suggests a benefit to estimating an independent WTP value for OD markets with dominant relative path quality (in addition to the network-wide estimate for PCAT2 and PCAT3 paths). Taking this idea to its limit, we imagine the best performing RM system can ultimately estimate passenger WTP in each of an airline's OD markets. While recent research efforts in PODS have begun to examine passenger demand at this most disaggregate level, so far these efforts have proved unsuccessful, with a primary obstacle being the scarcity and variability of booking data on an OD market basis. Nevertheless, we believe continued study of passenger WTP estimation can provide dividends not only in improving Hybrid Forecasting, but in future airline Revenue Management endeavors.