RM Methods for Airline Fare Family Structures

by

Vincent B. Surges

B.A., Mathematics, The College of St. Scholastica, 2011

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of

Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2013

Author . . . .
Department of Aeronautics and Astronautics
May 10, 2013

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

Accepted by . . . .
Eytan H. Modiano
Professor of Aeronautics and Astronautics
Chair, Graduate Program Committee

RM Methods for Airline Fare Family Structures

by

Vincent B. Surges

Submitted to the Department of Aeronautics and Astronautics
on May 10, 2013, in partial fulfillment of the
requirements for the degree of
Master of Science

Abstract

The rapid growth of low cost carriers forced many legacy airlines to simplify their fare structures and develop new pricing strategies to remain competitive. The strategy of branded fares, or "fare families", is an increasingly popular approach for airlines to differentiate their products and services from other competitors.

This thesis provides a comprehensive overview of revenue management (RM) forecasting and optimization methods developed specifically for fare family structures. These methods, collectively termed Q-Forecasting for Fare Families (QFF), provide airlines with the capability to manage branded fares from a RM perspective. The QFF methods are all constructed based on the assumed fare family passenger choice model, which accounts for both willingness-to-pay estimates as well as family preference. Each formulation makes underlying assumptions regarding passenger sell-up and buy-across.

The Passenger Origin Destination Simulator is used to test and compare the performance of each QFF formulation in a dual airline competitive environment, both with leg-based RM controls as well as network RM controls. The results from the simulations indicate that substantial gains in both revenue and yield over traditional RM methods can be achieved with appropriate RM in a fare family structure. Specifically, while Hybrid Forecasting (with leg RM controls) generates a 4.0% increase in revenue over Standard Forecasting, QFF is shown to increase revenues by more than 12.5%. The benefits of QFF are greater with network RM controls, with potential revenue increases of nearly 14.0% (over Standard Forecasting).

The positive results obtained with each QFF formulation are dependent upon an appropriate estimate for passenger sell-up and family preference. Consequently, this research also illustrates the importance of the estimate for passenger willingness-to-pay and its relationship to forecasting and optimization in airline RM.

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

Acknowledgments

First and foremost, I would like to express my sincere gratitude to my research advisor, Dr. Peter Belobaba. He introduced me to the field of revenue management, and made the transition to MIT extremely smooth. He continually provided guidance and support to me, in addition to always making himself available for any questions I might have had. This thesis would not be possible without his continual support; I am truly grateful for everything he has provided for me, least of the all the opportunities I have after MIT. His vast knowledge of the airline industry and professional intuition towards revenue management has contributed tremendously to my research findings.

I would also like to thank Craig Hopperstad, the developer of the Passenger Origin Destination Simulator (PODS), the simulation tool on which my results are based. He has provided extensive technical support on a wide variety of issues within PODS. His patience and guidance over the past two years, as well as his explanations of intricate concepts, all contributed to this thesis.

I extend my appreciation to all the members of the PODS Consortium: United Airlines, Delta Airlines, Boeing, Air Canada, LAN, Lufthansa, and SAS. I had the opportunity to participate in five PODS conferences during my two years at MIT. These conferences provided invaluable advice and support for my results, and provided me guidance on future research based on the members' experiences.

I am very grateful to have studied at MIT over the past two years. My time here has given me unparalleled research skills, while enhancing my skill sets in many other areas. I know that my tenure at MIT would not have been as exciting if not for all of the friendships that were formed both within my lab as well as with other graduate students throughout campus. I would like to especially thank Pierre-Olivier, Michael, Natasha, Tamas, Vishnu, and Alexander. Our time together was truly memorable, and I look forward to corresponding with each of you as we start the next chapter of our lives. And to the rest of the International Center for Air Transportation lab, including Eric, Mike, Karim, and Alyona, I want to thank you for making my time at MIT so enjoyable.

Most importantly, I must heartily thank my family and friends for their unconditional love and support. In particular, my wife Amanda and my parents. Without the constant support and understanding from my wife, these last two years would not have been possible. Amanda, I am eternally grateful for your love, and for you always being there for me, even when MIT tried to interfere! And to my parents, I am very grateful for everything you have done to make this moment possible. You both have always supported me and encouraged me to achieve my highest educational goals. Lastly, I would like to extend an appreciation out to the rest of my family for their constant support.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

After the Airline Deregulation Act passed in 1978, US airlines were granted the freedom to determine the price of their fares, routes flown, and frequency of domestic flights without government approval. Since then, the airline industry has evolved tremendously to adapt to an increasingly competitive environment. It was in the post-deregulation era where the practice of airline revenue management (RM), described by American Airlines as "selling the right seats to the right customers at the right prices" (Smith et al., 1992), began to develop. More formally, airline RM serves to design and manage service products by allocating seats on an aircraft to different predetermined booking classes in order to maximize total revenue (Weatherford, 1991).

While the nature of the airline industry has changed dramatically over the past three decades, the goal of maximizing revenue remains the same today as it did some 30 years ago. Many revenue management systems have been implemented over the years, owing to the advanced technological capabilities as well as the emergence of "low cost carriers" (LCCs) in the early 2000s. Many of the largest airlines, or "network legacy carriers" (NLCs), struggled to promote brand awareness in the early parts of the 21st century. One response to this issue was an innovative fare structure known as "fare families", first developed in 2004 in an attempt to de-commoditize the airline product.

The goal of this thesis is to provide a comprehensive overview of the different forecasting and optimization methods developed specifically for fare family structures. These models are analyzed using the Passenger Origin-Destination Simulator (PODS), a simulator used to test different airline RM systems in competitive RM scenarios.

1.1 Overview of the Airline Industry

The number of flights and fare options available to consumers wishing to travel from an origin to a destination (OD) increased tremendously after US airline deregulation. Airlines began to greatly expand the practice of differential pricing, an important component of revenue management (Botimer & Belobaba, 1999). Differential pricing consists of offering multiple combinations of price levels and restrictions, or "fare products", for a shared inventory of seats on the same flight leg. When an airline offered only one or two fare products, revenue was lost either in the form of consumer surplus or unused seats (Belobaba et al., 2009). The concept of differential pricing is illustrated in Figure 1-1. Higher revenues can be generated by offering multiple fare products based on passengers' expected willingness-to-pay (WTP). Each of the rectangular surfaces under the demand curve represent the revenue for a particular fare product. $Q1$ passengers book the fare product priced at $P1$, and so forth. As demonstrated by the figure, the airline's goal is to fill the empty space beneath the demand curve (lost revenue due to consumer surplus).

Although it is impossible in practice to charge each passenger their maximum WTP, airlines have made progress towards this goal through the use of differential pricing. When considering the different fare products an airline offers, consumers must make a trade-off between the restrictions associated with lower-priced fare products and the flexible (but higher-priced) unrestricted products. While some travelers find the multitude of fare product offerings overly-complex, economic theory supports the airlines' tactics.

In order to segment demand, airlines developed fare products that would appeal to different groups of passengers. Many legacy carriers offered a range of fare products, each with their own fares, or "price points", and set of restrictions. In addition to

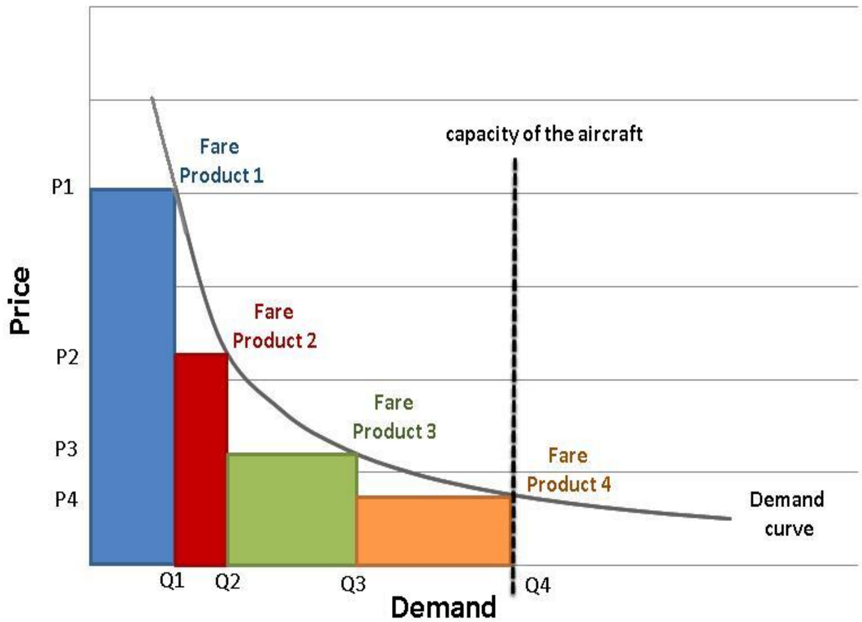

Figure 1-1: Airline Practicing Differential Pricing (d'Huart, 2010)

the different fare products offered, different cabins (economy, business, and first) were also offered, with multiple products within each of those cabins. This allowed those consumers who placed a higher value on travel flexibility (typically business travelers) the option to purchase a more expensive fare product in return for more convenient air travel. Alternatively, seats that were previously empty could be sold at lower price points to consumers who were more price-sensitive and less affected by restrictions (leisure travelers), thereby raising revenues (Belobaba et al., 2009). By recognizing the sensitivity of different types of passengers to the restrictions, the demand could be more accurately segmented.

The airlines imposed the most severe restrictions on lower fare products to prevent consumers with high WTP from purchasing these less-expensive products (a process known as "diversion"). Some of the restrictions included the non-refundability of a ticket, change fees, required Saturday night stay at a destination, as well as an advanced purchase (AP) requirement.

For RM purposes, every fare product is assigned to a booking class; there can be multiple products assigned to the same class. A collection of booking classes offered by an airline constitute a fare structure. An example of a typical restricted fare structure is shown in Table 1.1. The discounted M class is fully restricted and also requires passengers to book at least 21 days in advance. As the price points of each of the booking classes increases, the severity of the imposed restrictions generally decreases. The Y and B classes (top two booking classes) are attractive to business passengers, given that these types of passengers are generally unable to book many days in advance, and are unwilling to stay over Saturday night at their destination.

Booking Class	Fare Price	Advanced Purchase	Saturday Night Minimum Stay	Change Fees	Non-Refundable
Y	$900	0 days	No	No	No
B	$600	3 days	No	Yes	No
Q	$400	7 days	Yes	No	Yes
H	$300	14 days	Yes	Yes	Yes
M	$200	21 days	Yes	Yes	Yes

Table 1.1: Example of a Differentiated Fare Structure

Since passengers differ in their sensitivity to each of the restrictions, the costs associated with each of the disutilities is generally different for every traveler. When purchasing a fare product, a given passenger takes the restrictions into consideration by examining the total generalized cost of the fare product, that is, the sum of the fare itself along with the disutility costs the passenger in question attributes to the particular product. Out of the different options available, the passenger chooses the fare product with the lowest total generalized cost, provided the fare is below the passenger's maximum WTP.

The conventional applications of differential pricing and revenue management used in the 1990s were successful in limiting revenue dilution from diversion because of the restrictions associated with the cheaper fare products. With the legacy carriers focused on providing the most flexibility and gaining the most revenue from high-paying business travelers, the door opened for new entrant low cost carriers. These LCCs began to offer non-stop point-to-point service in major markets that had previously only been served by the NLCs. With the network carriers' pricing policies gradually leading to more expensive fares at the top, some business passengers avoided the higher prices by opting to stay over Saturday night in addition to meeting other restrictions imposed by the airlines. As a result, the legacy carriers faced major challenges as the number of passengers in the top booking classes dropped considerably with business travelers buying lower fare products, and in some cases choosing to fly with the LCC competitor.

The LCCs typically implemented more simplified fare structures compared to their NLC counterparts. These structures were typically composed of cheaper fare products with few or no restrictions attached to them; in some cases the price points acted as the only differentiator between the products. These fare structures were generally more appealing to passengers, both in terms of the lower prices and in the reduced complexity of restrictions. Consequently, traditional carriers in many cases abandoned the fully-restricted fare structure in order to stay competitive with the LCCs and maintain market share. Major restrictions were eliminated, shorter advance purchase requirements were implemented, and the fares were reduced. An example of a less-restricted fare structure is shown in Table 1.2. In addition to fewer restrictions attached to each of the booking classes, the less-restricted fare structure has lower fares (relative to the fully-restricted structure seen in Table 1.1) and shorter AP requirements for each of the classes.

Booking Class	Fare Price	Advanced Purchase	Saturday Night Minimum Stay	Change Fees	Non-Refundable
Y	$600	0 days	No	No	No
B	$450	0 days	No	No	Yes
Q	$350	3 days	No	Yes	No
H	$250	7 days	No	Yes	Yes
M	$150	10 days	No	Yes	Yes

Table 1.2: Example of a Less-Restricted Fare Structure with Lower Fares

In addition to the rise of the LCCs in the early 2000s, travel sites such as Orbitz and Expedia became much more prominent, enabling passengers to more easily compare the different fare products offered by the airlines for the same itinerary. This further contributed to the reduced market share for many of the legacy airlines, leading some major carriers into bankruptcy.

1.2 Branded Fare Families

With the emergence of the LCCs and low-fare search engines, the airline product was in many respects perceived to be a commodity. Many travelers selected an airline based on price and convenience rather than on any specific service benefits or loyalty to that airline. In response, some carriers attempted to develop new tools to both de-commoditize and enhance their product. One such strategy that was introduced by Air New Zealand (ANZ) in 2004 was branded fare families, a new approach to airline pricing and segmentation (Vinod & Moore, 2009). Since it was first developed, many other airlines have adopted this fare structure, including Air Canada and Qantas. This innovative fare structure consists of offering two or more sets, or "families", of booking classes, differentiating itself from the previous fare structure which consisted of a single set of multiple classes.

Each family in a fare family structure is set apart from one another by "fences" in service, flexibility, and fare levels. These differences exist between families, with the intention of providing the various passenger types (e.g., business, leisure) their own individual set of classes to choose from. *Within* each fare family all the classes have identical restrictions and amenities; the fare level of an individual class is the only differentiator between one class and another.

Figure 1-2 illustrates a fare family structure as offered by Air New Zealand for a particular South Pacific flight. For any itinerary, the cheapest available fare product within each of the four different economy families is presented as an option to the passenger at the time of booking. That is, ANZ's fare structure consists of four economy families, as well as additional families in the business cabin (not shown here). This type of fare structure presents a passenger with a multitude of options.

For example, the fare for the 7:00am flight departure ranges from $219 (in the "Seat" family) up to $369 ("Works Deluxe" family). The final decision on which family the passenger selects depends upon the services and flexibility the passenger desires, as well as the passenger's maximum WTP. Note that in some cases different flights within the same family have different prices. This is because the cheaper classes are not necessarily available for each of the departure times listed, and as a result, traveling at 9:30am is more expensive than at 7:00am in this example.

Figure 1-2: Fare Family Structure Offered by ANZ for Sydney-Auckland Market. Data source: AirNewZealand.com

This type of fare structure was developed for two primary reasons: incremental revenues and better brand awareness (Vinod & Moore, 2009). By offering different levels of service and flexibility, better passenger segmentation can lead to incremental ancillary revenues (Fiig et al., 2012). Additionally, better product recognition can potentially be achieved, given that passengers can secure all of their desired services and amenities at the time of booking. For instance, referring back to the ANZ fare family structure, a passenger booking in the "Seat + Bag" family avoids any of the inconveniences associated with checking a bag on the day of departure. From an airline's perspective, offering multiple families of classes ideally encourages passengers to shop for specific families by name, and thus creates a better brand image.

The ability to choose between multiple families can also be beneficial from a consumer perspective. Business passengers who have the ability to book in advance (yet still desire a fare product with flexibility and amenities) now have the opportunity to purchase a lower-priced fare product within their particular family of choice. Contrast this to the previous fare structure (with only a single set of classes), where even if a business traveler was able to meet the advanced purchase requirements, he was still forced to purchase a fare product in one of the top booking classes (at the highest fare levels) to ensure receiving the desired services and amenities.

In a similar manner, fare family structures can also be attractive to the typical leisure passenger. Price-sensitive travelers who are unable to meet one or more restrictions (such as an advanced purchase requirement) have the ability to purchase a ticket later in the booking period in one of the cheaper but more restricted families (e.g., "Seat" in ANZ's fare family structure). Although booking closer to departure will force passengers in this situation to purchase one of the higher-priced products, they now have the ability to do so in a cheaper family. This is generally much less expensive than purchasing a product in one of the top booking classes in a fare structure with only one set of classes. Because of the advantages to consumers just mentioned, airlines that offer fare family structures generally see higher load factors (LF), that is, higher occupancy rates on their flights.

While there are several benefits of implementing a fare family structure to both airlines and passengers, there are major RM challenges that must be taken into consideration. Virtually all RM systems were developed under the assumption of independent demand by fare product, that is, the notion that the demand for a particular product was independent of the availability of fare products in lower classes. Although this assumption simplified many RM models, it was clearly unrealistic. With the industry evolving and the new entrant airlines making their presence in the industry known, new forecasting and optimization methods were needed to ensure full benefit of the fare family approach.

1.3 Objectives of the Thesis

The objective of this thesis is to provide a detailed analysis of the different RM methods developed in this research for a fare family structure. Forecasting and optimization, which will both be described in greater detail in the subsequent chapters, play a critcal role in the success of an RM system, and thus in revenues being maximized. This thesis will present the specific algorithms developed in this research for a fare family environment. The robustness of the forecasting methods will be tested by examining their performance in different types of fare family structures. The limitations of the different developed models will also be discussed.

All simulations and quantitative evaluations will be performed using the Passenger Origin-Destination Simulator (PODS), first developed by Hopperstad in 1994 at the Boeing Company as an evolution from its predecessor, the Decision Window Model. The key components of the simulator, which can be adjusted to conform to a particular area of study, include the passenger choice model and the airline RM system.

1.4 Structure of the Thesis

This thesis is structured as follows. Chapter 2 presents an overview of the previous work done on airline revenue management, emphasizing the changes in the industry. The topics that will be discussed in this chapter include RM in the 1980s and 1990s, emergence of LCCs, and forecasting in an unrestricted environment.

Chapter 3 provides an in-depth description of the forecasting and optimization models developed for fare family structures. The assumed passenger decision process in a fare family environment is introduced, providing the motivation for each of the methods. The key differences in each of the formulations is highlighted.

The simulator used to test the different methods, the Passenger Origin Destination Simulator (PODS), as well as the simulation environments are discussed in Chapter 4.

In Chapter 5, the performance of the different forecasting and optimization methods developed in Chapter 3 are analyzed. Specifically, the simulation results from the different test cases will be presented and analyzed using different metrics.

The thesis is summarized in totality in Chapter 6. The impacts of fare families are recapped, as well as potential future research directions in this field.

Chapter 2

Literature Review

Although airline RM literature dates back to the early 1970s, the field truly evolved in the post-deregulation era. With the goal of achieving the highest possible revenues, focus shifted from maximizing the quantity of passengers carried to optimizing seat allocation. This chapter will cover several of the most important published works in this area of research, beginning with an overview of the traditional RM methods developed under the assumption of independent booking class demand. The second section describes different forecasting techniques that were developed for less-differentiated fare structures, that is, methods that do not rely on the independent demand assumption. All of the methods discussed in this chapter were developed under the assumption of a fare structure with a single set of classes (in contrast to the fare family structures introduced in Chapter 1).

2.1 Traditional Airline RM Methods

In the 1980s, airline RM systems were first developed and used as large database management systems. Over the years more sophisticated features were added; the current third generation RM system implemented by many airlines throughout the world has capabilities to both forecast demand for each future flight departure and to optimize the number of seats made available to each booking class. It is the function of the RM system to set booking limits (BL) on the lower classes to protect seats for consumers who are willing to pay more for the less restricted product (Botimer & Belobaba, 1999). On the other hand, if the plane is not full, the RM system should ensure that empty seats be made available for passengers with a lower willingness-to-pay who otherwise would not fly. Weatherford (1991) provides more insight on the revenue opportunity which is lost whenever a flight departs with empty seats.

The major components of a typical third generation RM system are illustrated in Figure 2-1. Historical booking data from the same flight in the past (same route, time, and day of the week) is combined with the actual booking data to generate a demand forecast by booking class for the upcoming flight. This forecasting model, along with the estimated revenue data, are inputs into the optimization model that works to create booking limits on each booking class that will maximize revenue on the flight leg. During this time, an overbooking model is also used to estimate an appropriate number of seats of each class to make available given the historical no-show rates. Overall, the optimization model and the overbooking model create the recommended booking limits per class for the flight leg in question.

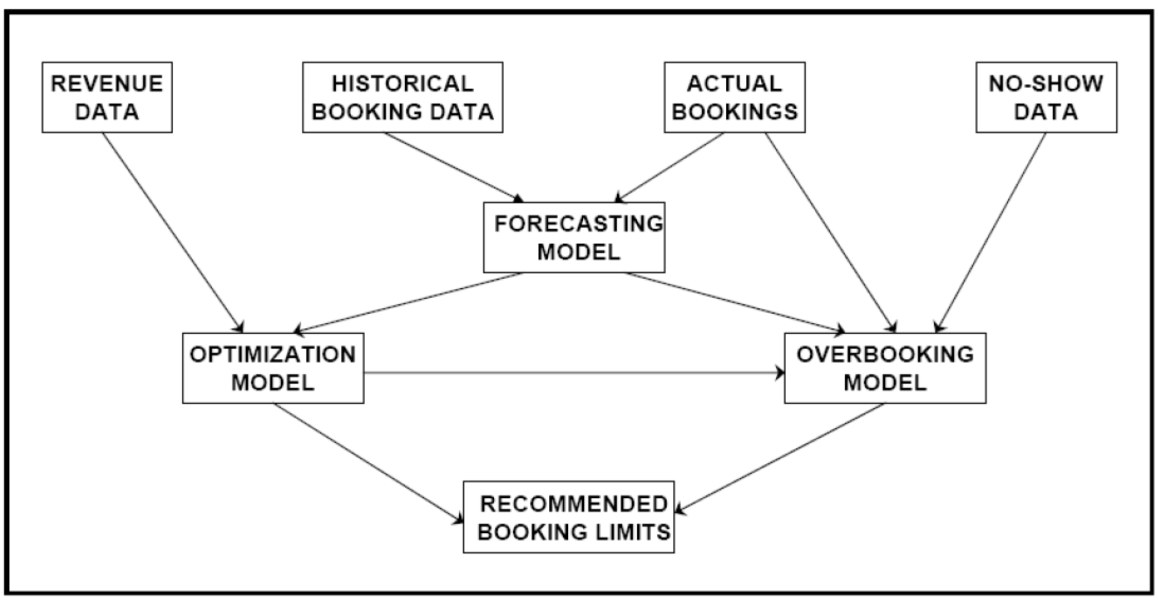

Figure 2-1: A Typical Third Generation RM System (Barnhart et al., 2003)

Effective implementations of standard RM systems have been estimated to increase revenues by 4-6% over situations in which no seat inventory is used (Belobaba, 1992b). This chapter will discuss two of the three main components to the third generation RM system: optimization models and forecasting. For a detailed overview of the history of RM in the airline industry, consult McGill & Ryzin (1999) and Barnhart et al. (2003).

2.1.1 Seat Allocation Models

One of the main complications in airline RM is that low-yielding leisure passengers generally book much earlier in the booking process than high-paying business travelers. This creates a need for inventory control based on forecasts of various passenger types. That is, a sufficient number of seats should be protected for high-fare passengers who book closer to departure.

2.1.1.1 Leg-based Controls

Littlewood (1972) developed the first model for seat inventory control by solving the booking class problem for two 'nested' classes (more on nested classes below). Belobaba (1987 and 1989) developed the Expected Marginal Seat Revenue (EMSR) heuristic, a more general model that is applicable to any number of nested booking classes. In 1992, Belobaba subsequently developed a variant (EMSRb) of this algorithm to make it more robust. EMSRb has been widely incorporated into many airlines' RM systems on the flight leg level and is used extensively throughout this thesis.

The EMSRb model assumes demand to be stochastic (normal) and independent for each booking class, and assumes that the lowest classes book first. The model determines the leg-based nested booking limits, that is, protection levels for higher classes and booking limits on the lower classes. In general, airline reservation systems use a nested control mechanism, preventing higher-priced booking class requests from ever being denied so long as the plane is not full. Figure 2-2 shows the general idea of nested booking limits, where there are booking limits on each of the four classes (Y class being the most expensive and the Q class being the cheapest). The booking limit for the Y class, or BL_1, is set equal to the remaining capacity of the plane.

In the EMSRb model, a seat is protected for the higher booking class whenever the revenue expected from protecting it exceeds the revenue from the fare directly below it. Belobaba & Weatherford (1996) provide a much more in-depth explanation of the EMSRb heuristic.

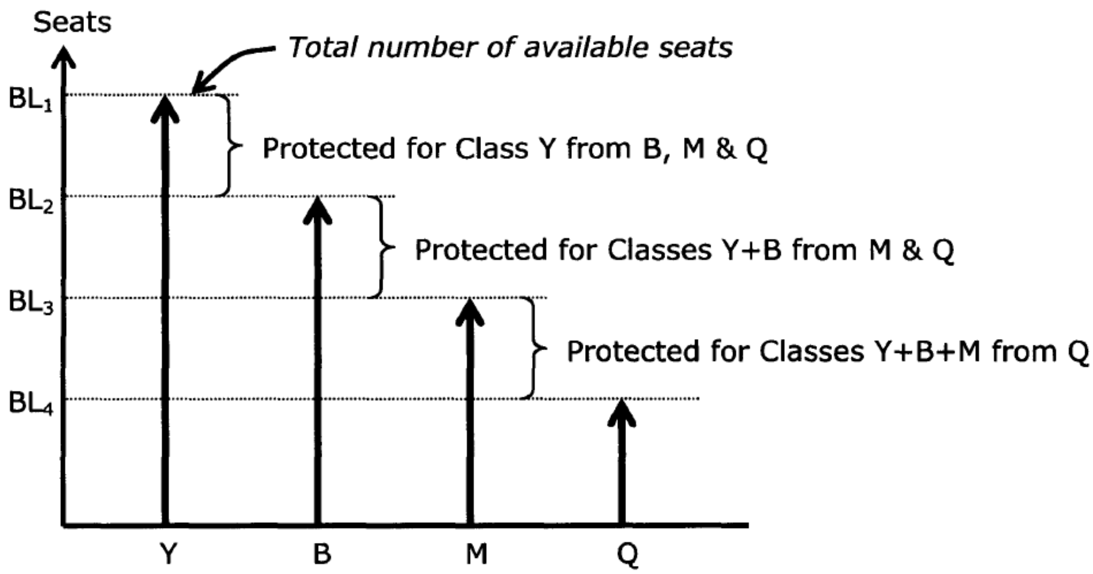

Figure 2-2: Nested Bookings Limits and Class Protection Levels (Cleaz-Savoyen, 2005)

The EMSRb model optimizes seat protections levels on a single leg basis. However, it does not take into consideration the difference in revenue from local (single-leg) vs. connecting (multiple-leg) passengers. Belobaba et al. (2009) identifies two potential problems that prevent a leg-based optimization approach from maximizing revenue. The first occurs when a short-haul flight leg is full, acting as a bottleneck for other itineraries connecting multiple legs. Revenue is lost when high-revenue passengers are unable to purchase their desired itineraries due to no available seats on one of the flight legs on their journey. The second issue with leg-based control is the inability to distinguish between passenger origin-destination itineraries. When a connecting itinerary is accepted by the airline, there is potential that this one passenger can prevent two high-paying local travelers from flying. Given that the sum of two local itineraries is generally greater than one connecting itinerary, revenue is lost when this occurs. For these reasons, leg-based control is sub-optimal because it favors local passengers. RM methods which take into account these network effects, classified as Origin-Destination (OD) controls, are discussed next.

2.1.1.2 Origin-Destination Controls

OD controls are especially important given that the majority of major airlines operate hub-and-spoke networks (de Neufville & Odoni, 2002), and thus will have many multi-leg itineraries. There has been substantial effort in developing algorithms for origin-destination, or path-based, control of booking classes. One of the first approaches to network revenue management (i.e., RM at the path level) was developed by Smith & Penn at American Airlines in 1988. In this approach, booking classes are replaced by "revenue value buckets" for seat inventory management. These buckets are defined according to network revenue value. Each origin-destination itinerary fare (ODIF) in the network is assigned to a revenue value bucket. The seat availability for a requested ODIF (local or connecting) depends on the availability of the corresponding revenue value bucket for each leg of the passenger's itinerary. With this approach, different itineraries can be compared to one another in terms of total revenue value. The pitfall of relying only on revenue value is that it is "greedy" in the sense that it always favors connecting passengers, even in situations where two locals would contribute more overall revenue.

The greedy virtual bucket method was refined into Displacement Adjusted Virtual Nesting (DAVN). In this new methodology, the revenue input to the virtual buckets was adjusted for the network displacement costs incurred by a connecting passenger, given that the actual network contribution of a connecting passenger is less than or equal to the total ODIF value of the passenger's itinerary (Belobaba et al., 2009). Williamson (1992) describes different methods that can be used to estimate these displacement costs. After incorporating the network displacement costs, leg-based controls can then be applied to solve the problem. A good overview of these models and their evolution can be found in Belobaba (2002). Implementing network RM controls can lead to an additional 1-2% increase above leg-based controls (Belobaba, 1998). This thesis uses EMSRb and DAVN as the main seat allocation models in the simulations involving fare family structures.

2.1.2 Demand Forecasting Models

Another component of the RM system that that plays a vital role in maximizing revenue is the demand forecasting model. The objective of forecasting is to estimate future values for a particular flight leg, including demand for a booking class as well as passengers' willingness to "sell up", that is, purchase a higher fare product if their first choice is no longer available. The forecasts must be as accurate as possible for the RM systems to be effective. It is common that the observed bookings are less than the actual demand; this happens when a particular booking class is closed and passengers are rejected for that particular class. When this occurs, the historical bookings will not reflect true demand, but rather will represent a lower bound on the actual demand. Hence, it is crucial that the booking data be detruncated to estimate the unconstrained demand. For an overview of some of the unconstraining methods used, please refer to Weatherford & Polt (2002). The most common and frequently used forecasting methods are now discussed.

2.1.2.1 Pick-up Forecasting

Pick-up Forecasting is a simple forecasting technique that has proven to be effective under the traditional assumptions of RM, specifically, the assumed independence of demand in different booking classes. This forecasting method calculates the expected incremental bookings for each data check point (or "time frame") until departure using the historical database. This pick-up is added to the number of bookings already received to forecast the total demand at the end of a specific time frame. There are two different versions of this model: the classical and the advanced pickup model. The classical model only uses data from flights that have departed, while the advanced pick-up model developed by L'Heureux (1986) also uses data from flights that have not yet departed. In this thesis, only the classical pick-up forecasting method is used in the simulations. Gorin (2000) provides a detailed description of this forecasting algorithm.

2.1.2.2 Alternative Forecasting Methods

A slight variant to the traditional Pick-up Forecasting is Pick-up Forecasting with Exponential Smoothing. There is one additional parameter in this method which allows the user to put more weight on recent samples as opposed to the older ones.

Another classical forecasting method is Regression Forecasting. This method assumes a correlation between bookings-on-hand and future bookings for a particular flight. The forecasted demand thus directly depends on the bookings-on-hand at each specific time frame prior to departure. For a detailed explanation on Regression Forecasting, refer to Zickus (1998). Consult Weatherford (1999) for an overview of common forecasting methods used in practice.

2.2 RM Methods for Less-Restricted Fare Structures

In the late 1990s, most legacy airlines were upgrading their revenue management systems to include network control, unaware of the major changes that would be taking place after the turn of the century. It was at this time when the LCCs started to emerge. With the low-fare search engines freely available on the internet, passengers had easy access to information on the different fare products offered by the airlines. To remain competitive, many of the legacy airlines underwent dramatic changes to their pricing practice and adapted to the LCC fare structure, removing many of the restrictions which acted as fences between the different products. Assumptions such as fare product demand independence, crucial to the success of standard forecasting and optimization methods, became completely invalid.

Since every passenger buys the lowest fare product when no restrictions are attached, it was not possible to accurately forecast future demand simply based on historical data. As a direct result, the traditional revenue management systems

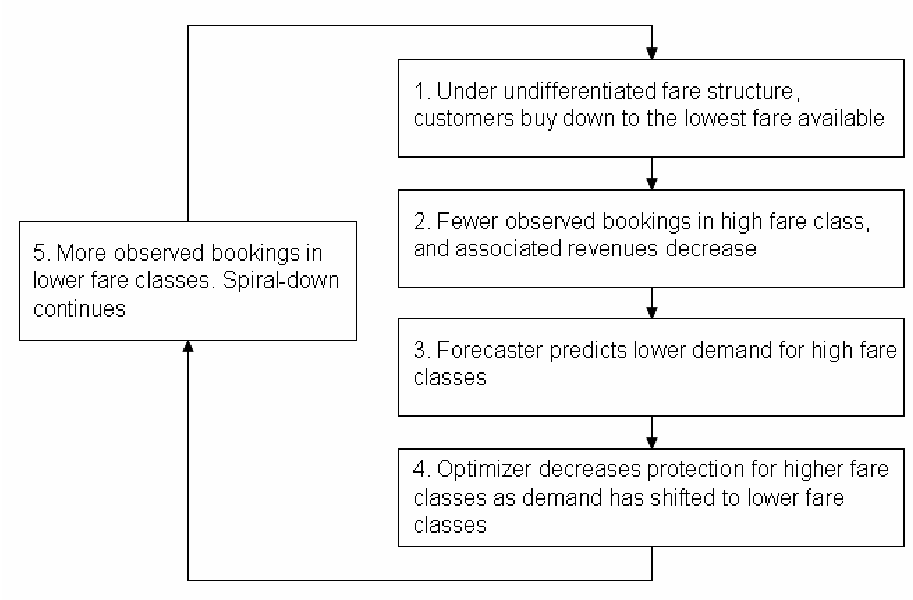

Figure 2-3: Spiral-Down Effect (Tam, 2008)

failed, both in forecasting and in optimization. Demand for higher booking classes was under-forecasted, which then led to fewer bookings in higher classes and more bookings in lower classes, and thus tremendous revenue losses. This cycle of diversion, referred to as "spiral-down", is shown in Figure 2-3. For more on the effects of spiral-down and less restricted fare structures, consult Cooper et al. (2006) and Cleaz-Savoyen (2005).

To avoid spiral down, airlines initially applied manual interventions to override the RM system's recommended booking limits. However, this was not the long-term solution. For the NLCs to return to higher revenues, they were forced to either modify their current RM systems or develop new tools that would allow them to capture passengers with high WTP. Widespread research has been done on forecasting passenger demand for unrestricted fare structures. The remainder of this chapter will focus on methods that have been developed specifically for the these types of fare structures. The central idea behind many of these methods is estimating passengers' WTP, and then forecasting based on this concept, rather than relying on the demand independence assumption and time series forecasts.

2.2.1 Q-Forecasting

Q-Forecasting was developed by Belobaba & Hopperstad (2004) as a forecasting method for fully undifferentiated booking classes that was designed to avoid the spiral-down effect. Q-Forecasting does not use the independent demand assumption, but rather only forecasts potential demand for the lowest available booking class, given this class contains the only fare product passengers will purchase when the booking classes are undifferentiated.

The Q-Forecasting process is shown in Figure 2-4. The first step is to obtain unconstrained historical demand by class (for each of the remaining time frames), and then convert this demand to the total number of "Q-equivalent bookings", that is, the equivalent demand for the lowest booking class if it were available. To accomplish this, sell-up probabilities between the lowest booking class (Q) and the rest of the booking classes must be established. The probability of sell-up from class Q to some higher booking class f depends on the price points of both of the classes, as well as the airlines estimate of passengers' WTP.

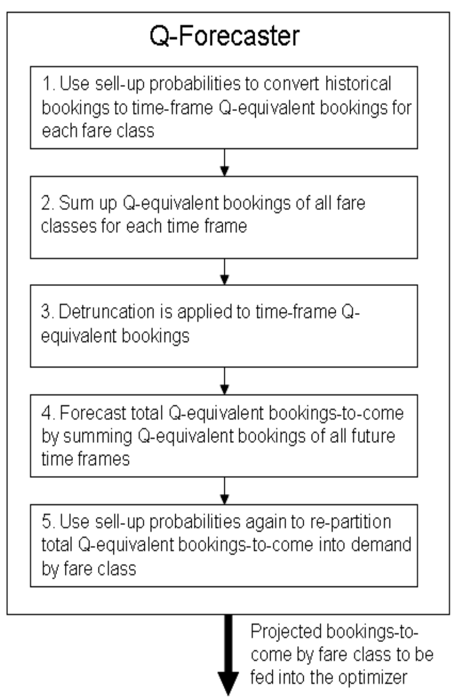

Figure 2-4: Process of Q-Forecasting (Belobaba, 2010)

In the next step, detruncation is applied to the Q-equivalent bookings for each time frame. These unconstrained Q-equivalent bookings are then summed up to produce a total Q-equivalent bookings-to-come within the current time frame. The sell-up probabilities from class Q to the higher classes are then used to partition the Q-equivalent bookings into the different booking classes to develop the total bookings-to-come forecast by class. Note that the estimated sell-up probabilities will change over time prior to departure, given that late-booking travelers are generally less price-elastic and will be more likely to pay extra for a ticket than passengers booking months in advance. To account for this, a weighted average of the sell-up probability for each class is used. Given that sell-up can only be achieved by closing down lower booking classes based upon the estimated passengers WTP, appropriate sell-up probabilities are crucial to the effectiveness of Q-Forecasting. Cleaz-Savoyen (2005) discusses various methods for estimating sell-up probabilities. In this thesis, sell-up is modeled as a negative exponential distribution (Belobaba & Hopperstad, 2004).

2.2.2 Hybrid Forecasting

The Q-Forecasting method introduced in the last section was developed for fully unrestricted fare structures. However, most airlines' (including LCCs) fare structures still have some restrictions, given that fare restrictions are generally only partially removed. Even with this type of structure, differentiating between business and leisure passengers was much more difficult than it was with a fully-restricted fare structure. Boyd & Kallesen (2004) suggest that the demand can instead be differentiated between price-oriented and product-oriented passengers. Price-oriented passengers are those passengers who always purchase in the lowest open booking class, regardless of what other options are available. Passengers are considered to be product-oriented if they purchase a fare product in any class above the lowest open one. That is, they purchase a higher-priced product because of its flexibility and/or services.

By segmenting the demand in this manner, different forecasting methods can be applied to the two different "groups" of passengers. Q-Forecasting can be applied to the price-oriented passengers while standard Pick-Up Forecasting can be implemented for the product-oriented passengers. Hybrid Forecasting is thus a combination of the two methods that forecasts the demands separately and then aggregates the two forecasts to estimate total forecasts. Figure 2-5 shows a schematic of the Hybrid Forecasting process. Reyes (2006) provides more information on the Hybrid Forecasting methodology and its performance

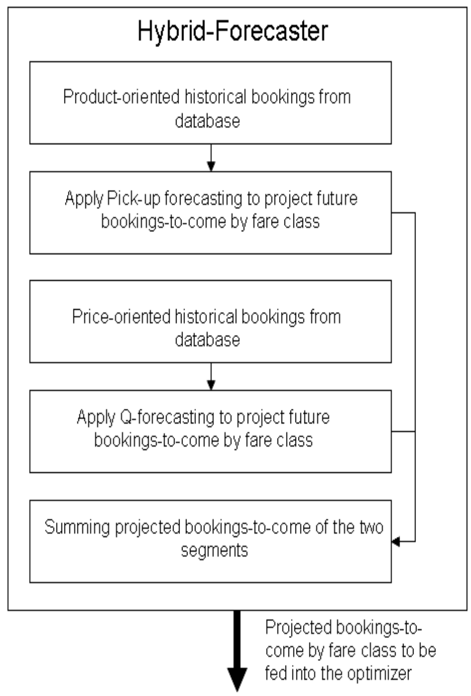

Figure 2-5: Process of Hybrid Forecasting (Belobaba, 2010)

2.2.3 Fare Adjustment Theory

Both Q-Forecasting and Hybrid Forecasting were efforts to reversing the spiral-down effect. However, it was the fare adjustment theory developed by Fiig et al. (2010) that provided the optimization models for the forecasting methods mentioned above. The fare adjustment theory showed that it was possible to map the fare and demand inputs of a completely general fare structure (perhaps unrestricted) into an equivalent independent demand model of a fully restricted structure. This is known as the marginal revenue transformation theorem. Once the demands and fares have been transformed, the traditional RM methods that were developed under the independent demand assumption can then be applied using the original fare structure (Fiig et al., 2010).

In a fully restricted fare structure with price points in decreasing fare order and deterministic demand for each fare product, the demand for each fare product is assumed to be independent of whether other products are available. It is therefore beneficial to open up classes in decreasing fare order until capacity is reached, given that total revenue increases monotonically with the total quantity of seats sold.

In an unrestricted fare structure in which the price points act as the only differentiator between the classes, passengers who had previously been willing to purchase the more expensive fare products will now buy in the lowest available booking class. Opening up classes in decreasing fare order will result in passengers diverting to the lowest products. While the incremental demand from opening up the next lowest class remains identical to the incremental demand in a fully restricted fare structure, the incremental revenue is reduced, taking into account the fact that passengers are diverting to lower-priced classes.

The optimization problem is thus to maximize total revenue subject to the quantity sold being less than the capacity of the flight leg. This solution is most easily arrived at by considering the total forecasted demand and revenue for all possible combinations of classes, or "policies", an airline can offer (which is 2^n in an n class structure). As Figure 2-6 illustrates, the policies trace out the convex hull, with the optimal policies $S_0...S_m$ identified as those that lie along the 'efficient frontier' (upper boundary) on the convex hull (Fiig et al., 2010). As illustrated by the figure, the marginal revenue f_i', or "adjusted fare", is the change in total revenue ($TR_i - TR_{i-1}$) over the change in total demand ($d_i' = Q_i - Q_{i-1}$) from one policy to another.

Assuming a nested efficient frontier, the marginal revenue transformation will produce demands and adjusted fares for each of the original booking classes. After these quantities are fed into the optimizer, the same availability control as in the independent fare structure can be achieved.

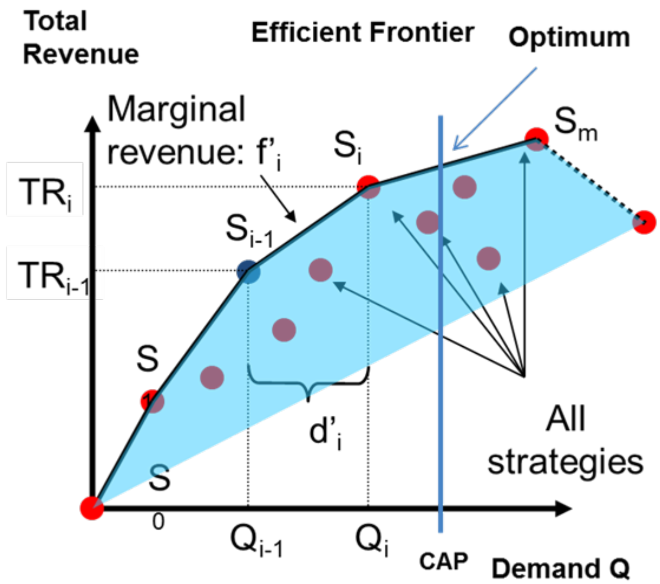

Figure 2-6: Scatter Plot of Different Strategies Tracing out the Convex Hull (Fiig et al., 2010)

2.3 Chapter Summary

This chapter began with a review of the literature on airline RM methods that were in use prior to the emergence of the LCCs. The discussion was centered on both seat allocation optimization and forecasting models. Next, the literature focused on revenue management methods in a less-restricted environment. The impact of LCCs on legacy carriers has been immense, as the removal of fare class restrictions and cheaper fares has led to lower revenues for the NLCs. Finally, the fare adjustment theory was introduced. The next chapter focuses on RM methods developed in this research for fare family structures.

Chapter 3

Revenue Management Methods for Fare Families

One of the main challenges in implementing a branded fare family structure in practice is that it requires more sophisticated forecasting and optimization models. That is, the RM methods discussed thus far were not developed for a multiple-family fare structure. Although implementing the traditional RM methods (e.g., EMSRb optimization with Hybrid Forecasting) in a fare family environment may provide a small gain over manual control, there is an additional revenue gain that can be earned through successful modeling of customer choice behavior in the RM system. Given that some experts believe that the fare family structure could be the key to effective airline pricing in the future, developing RM models specifically for this environment is essential (Fiig et al., 2012). This chapter introduces the advanced forecasting methods developed for the branded fare family structure, providing an in-depth overview of each algorithm.

3.1 Introduction to Fare Family Structures

The fare family structures used for examples and simulations in this thesis are comprised of two families: 'family 1' (known as the flex fare family in some literature), and 'family 2' (restricted fare family). It is worth mentioning that in practice an airline can offer any number of families. The families are differentiated from one another by both the services offered as well as flexibility. Within each family the booking classes are undifferentiated in terms of restrictions; the lower-priced family 2 classes are completely restricted while the more expensive family 1 classes are completely unrestricted. Thus the only difference between classes in each family is the price points, which are in decreasing fare order.

In this thesis, the cheapest available (or "open") family 1 and family 2 classes are denoted as $f1$ and $f2$, respectively. An optimization policy $\{f1, f2\}$ consists of opening up at most one class from family 1 and at most one class from family 2. It suffices to consider only the price point of the lowest open class in each family, since passengers will always book in the cheapest class when choosing between options in the same family with identical restrictions. Thus in any fare family structure with n family 1 classes and m family 2 classes, the maximum number of possible policies to be considered for forecasting purposes is $(n+1)(m+1)$, which also accounts for the policies where one or both of the families are completely closed.

3.2 Passenger Decision Process

In order to develop successful forecasting methods for fare family structures, the passenger decision process must be modeled appropriately. Modeling this process is much more complicated in a fare family structure than in a fare structure with a single set of booking classes (referred to as a 'nested fare structure' in this thesis, given that the majority of non-fare family fare structures have the nested property). In a nested fare structure, a passenger simply chooses a product in the cheapest class that meets his criteria, or is a "no-go", that is, the passenger does not fly. In a fare family structure, a passenger now has the option to choose between the lowest open family 1 class $f1$, the lowest open family 2 class $f2$, or neither.

Before proceeding further, formal definitions of two important concepts in fare family structures are provided. "Sell-up", briefly described in Chapter 2, refers to the decision of a passenger to purchase a higher-priced fare with the same (or very similar) restrictions if the passenger's first choice of booking class is unavailable. The decision to sell up is based strictly on passengers' WTP, and only passengers with a WTP greater than or equal to the higher-priced fare will sell up. Given that there is no difference in restrictions between the booking classes within each family, the notion of passengers' willingness to sell-up plays a critical role in the forecasting techniques developed throughout this chapter.

Another important concept is the notion of "buy-across". Buy-across refers to the choice an eligible passenger (i.e., a passenger that can afford both $f1$ and $f2$) makes between the lowest open family 1 class $f1$ and the lowest open family 2 class $f2$. This choice is made by comparing the difference in price points between $f1$ and $f2$ against the difference in disutilities associated with the restrictions and/or amenities of the two families. Given that family 1 is unrestricted in the examples used in this thesis, this is equivalent to comparing the difference in price points between $f1$ and $f2$ against the disutility costs the passenger in question attributes to the family 2 restrictions.

Figure 3-1 shows how sell-up and buy-across are incorporated into the assumed choice process in a fare family structure from a passenger's perspective. In this model, the first step is for the passenger to determine if he can afford $f2$, that is, determine if his maximum WTP is greater than the price of the fare product in the lowest open family 2 class. If he can not afford $f2$, he becomes a no-go. If the passenger can afford $f2$, he would then compare his maximum WTP to $f1$ to see if he can afford the lowest open family 1 class. If the passenger can afford $f2$ but not $f1$, he purchases $f2$.

If the passenger can afford $f1$, he then compares the disutility costs that he attributes to the $f2$ restrictions against the difference in price points between the two classes, $fare_{f1} - fare_{f2}$, to determine which class has a lower total generalized cost (sum of the price point and disutility costs), and thus which class the passenger will book in.

New forecasting and optimization methods are critical to the success of fare families if the actual passenger decision process for a fare family structure follows the

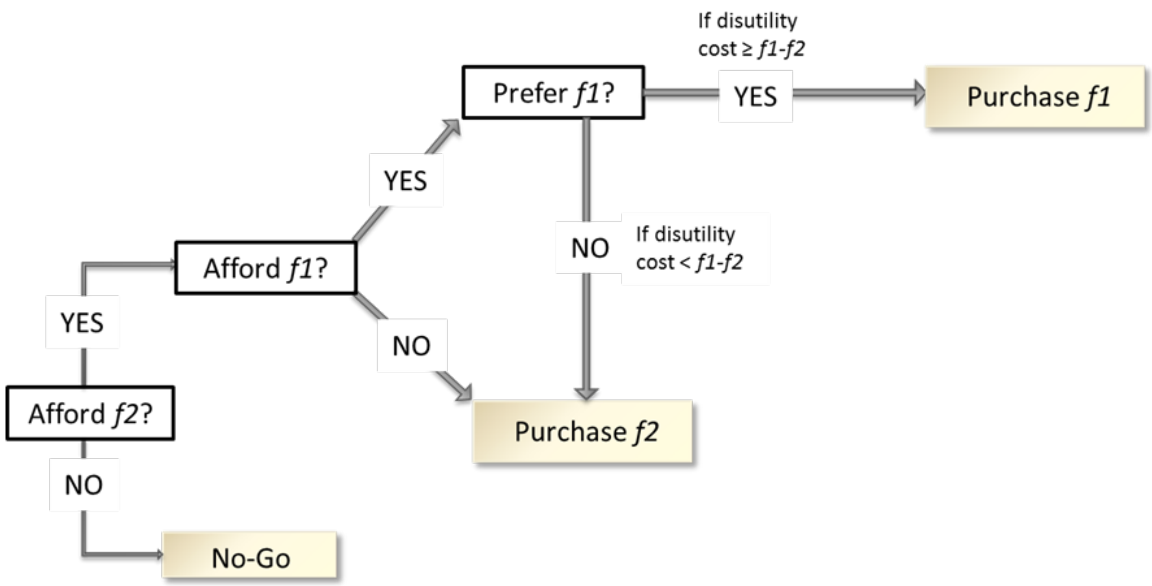

Figure 3-1: Assumed Passenger Decision Process in a Fare Family Structure

model described in Figure 3-1. That is, the forecasting methods described in Chapter 2 do not reflect the assumed passenger behavior in a fare family environment. Forecasting and optimization methods are needed that incorporate an airlines' estimate of passengers' WTP (sell-up), and that model passenger preference (buy-across) for every combination of family 1 and family 2 classes.

3.3 Q-Forecasting for Fare Families

The forecasting methods developed in this research for the fare family environment utilize the concept of Q-Forecasting (as described in Chapter 2) extensively. As such, the methods introduced in this chapter are collectively termed "Q-Forecasting for Fare Families" (QFF). Each of the methods described in this chapter follow a similar step-by-step process. These steps are all described in detail for the first method (QFF1). For each subsequently developed model, a similar procedure is followed. As such, when describing QFF2 and QFF3, only the steps in each method that differ from QFF1 will be highlighted. Note that, for each forecasting method, the RM system re-optimizes the booking limits at each time frame (TF) before departure.

To more easily illustrate the QFF algorithm, a numerical example highlighting each of the different steps in the QFF process (with the QFF1 formulation) is worked out. In this example, a two-family six-class fare structure is used (three classes per family). The classes in family 1 are labeled alphabetically from $A1$ to $C1$, with $A1$ being the highest-priced class overall and $C1$ being the cheapest family 1 class. Likewise, the family 2 classes are labeled in a similar manner from $A2$ to $C2$.

The fare structure used for this example is shown in Table 3.1. In this fare family structure, the price points are non-overlapping; the fare of the lowest family 1 class (in this case, $C1$) is priced above the most expensive family 2 class ($A2$). As indicated in the table, there is no advanced purchase requirement in either family. This will be the case for all simulations involving QFF; no AP is used in conjunction with any fare family forecasting method in this thesis.

Family	Booking Class	Fare Price	Restricted?	Advanced Purchase?
1	A1	$400	No	No
	B1	$300	No	No
	C1	$200	No	No
2	A2	$150	Yes	No
	B2	$125	Yes	No
	C2	$100	Yes	No

Table 3.1: Fare Family Structure used in QFF1 Example

3.3.1 QFF1

A flow chart summarizing the main steps in the QFF1 process is shown in Figure 3-2. After obtaining a Q-Forecast using historical data, the demand and revenue is forecasted for the different policies. Once the optimal policies have been determined, these policy forecasts can then be converted back into class forecasts, which can then be fed into the optimizer. A walk-through of QFF1 is provided in the following subsection, illustrating the process within a particular time frame. As mentioned above, these same steps are performed within each time frame in the pre-departure process.

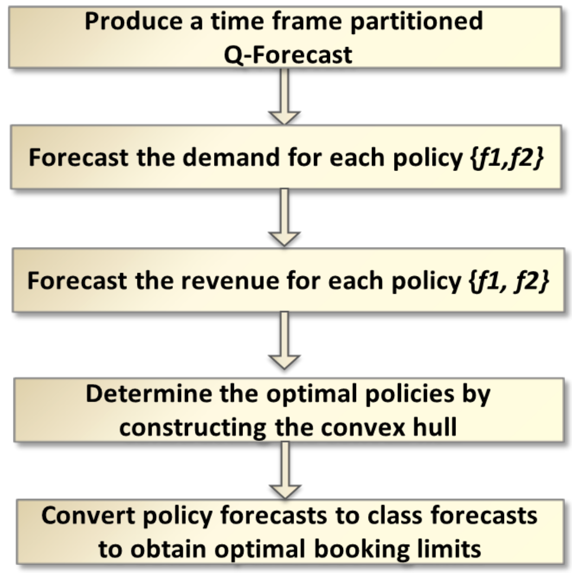

Figure 3-2: QFF1 Process

Step 1: Produce a time frame partitioned Q-Forecast

The first step in the forecasting process is to use unconstrained historical data to generate a time frame partitioned Q-Forecast both for the current time frame as well as for each future time frame up until departure. In the numerical example, the partitioned Q-Forecast for the current time frame tf, denoted as FC_{tf}, is assumed to be 30. Note that with QFF1, the Q-Forecast is generated assuming that $C2$ is the lowest class. That is, no distinction is made between the two families.

Step 2: Forecast the demand for each policy $\{f1, f2\}$

In this step, the maximum demand is forecasted for each policy. The maximum demand is estimated assuming that the family 1 and family 2 classes that make up the policy are the lowest open classes in each of their respective families. Given that the fare structure used in this example consists of three classes in each family, there are $(3+1)(3+1) = 16$ possible policies the airline can choose to offer. However, in a non-overlapping fare structure, an airline would never offer a policy consisting of only a family 2 class. Thus the actual number of policies that need to be considered in this setup is $(3)(3+1) = 12$, which corresponds to the nine policies with both family 1 and family 2 classes, as well as the three policies with only a family 1 class.

The demand forecast for each policy depends on the airline's estimation of passengers' willingness to sell up. A negative exponential distribution sell-up model is assumed in this thesis. This model is governed by a single parameter called "FRAT5", defined as the fare ratio of a higher fare to the lowest fare class ($C2$) at which the airline expects 50% of the demand will sell up to a higher class (Belobaba & Hopperstad, 2004). The probability a random passenger will sell up to class f (at a fare of $fare_f$) in time frame tf, given he would have chosen the lowest class $C2$ and no other classes are available, is

$$psup_{tf}(\text{C2} \rightarrow f) = e^{(\frac{fare_f}{fare_{C2}}-1)(\frac{ln(0.5)}{FRAT5_{tf}-1})}$$

One of the fundamental assumptions with QFF1 is that all passengers in the forecast are assumed to follow the same sell-up model. That is, the QFF1 algorithm does not distinguish between passenger types (e.g., business vs. leisure).

For a given flight, it is expected that as the departure date nears, passengers' WTP, and consequently the sell-up probabilities increase, given that less price sensitive business travelers tend to book closer to departure. To account for this, the single FRAT5 curve is modeled in such a way to capture the different booking tendencies between business and leisure passengers. Specifically, the sell-up probability is expected to gradually increase from the early time frames up until departure, resulting in a FRAT5 curve of an S-shape, as shown in Figure 3-3. This curve reflects the change in the business/leisure mix across time frames (Belobaba, 2010).

Forecasting the demand for any policy with QFF1 is dependent upon the lowest overall class in the policy. Consider any of the 9 policies where some family 2 class is open. The maximum demand for any policy of this type in time frame tf is the Q-Forecast in time frame tf multiplied by the probability of sell-up from $C2$ to $f2$ in time frame tf. This corresponds to the number of passengers from the Q-Forecast

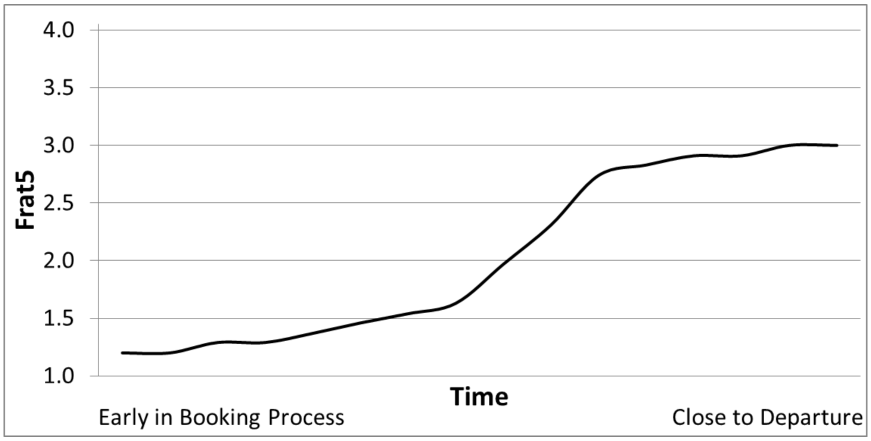

Figure 3-3: Typical FRAT5 Curve

that can afford $f2$:

$$dem_{tf}\{f1, f2\} = FC_{tf} \times psup_{tf}(C2 \rightarrow f2)$$

The maximum demand for any policy with $f2$ open is thus independent of the lowest open family 1 class. Note that this is always true for a non-overlapping fare family structure, since in this setup any passenger than can afford $f1$ can also afford $f2$. Thus only those passengers who can afford $f2$ need to be considered when some family 2 class is open.

Now consider the scenario where family 2 is completely closed (denoted with a %). With only the lowest open family 1 class $f1$ open, the maximum demand for $\{f1, \%\}$, that is, the estimated number of passengers from the Q-Forecast who can afford $f1$ in time frame tf, is found as follows:

$$dem_{tf}\{f1, \%\} = FC_{tf} \times psup_{tf}(C2 \rightarrow f1)$$

The previous two demand equations illustrate that the demand forecast for any policy with QFF1 only depends upon the cheapest available class in the given policy. If a family 2 class is open, the maximum demand for $f2$ will be the overall policy demand. If no family 2 classes are available, then the demand of the policy will be the demand for the lowest available family 1 class.

A numerical example of forecasting the maximum demand for the different policies is shown in Table 3.2. The calculations performed assume $FC_{tf} = 30$ (step 1), and a FRAT5 value of 2.0 in time frame tf. As the example illustrates, the demand forecast is the same for the policies $\{A1, A2\}$ and $\{C1, A2\}$, given that $A2$ is the lowest open family 2 class in each policy.

Policy	*f1*	*f2*	Max Demand
(A1,%)	A1	%	3.75
(A1,A2)	A1	A2	21.21
(A1,B2)	A1	B2	25.23
(A1,C2)	A1	C2	30.00
(B1,%)	B1	%	10.61
(B1,A2)	B1	A2	21.21
(B1,B2)	B1	B2	25.23
(B1,C2)	B1	C2	30.00
(C1,%)	C1	%	15.00
(C1,A2)	C1	A2	21.21
(C1,B2)	C1	B2	25.23
(C1,C2)	C1	C2	30.00

Table 3.2: Example Illustrating the Demand Forecast for each Policy assuming $FRAT5_{tf} = 2.0$

Step 3: Forecast the revenue for each policy $\{f1, f2\}$

In the assumed fare family passenger decision process the notion of sell-up and buy-across were both discussed. In the demand forecast, sell up from $C2$ to f in time frame tf was taken into consideration by the variable $psup_{tf}(C2 \rightarrow f)$. The demand forecasts for each policy corresponded to the maximum number of passengers from the Q-Forecast that could afford the lowest open class in the given policy. The revenue forecasts, discussed in detail in this section, also account for buy-across, given that some passengers in the demand forecast will purchase the family 1 product.

Similar to the demand forecasts, forecasting the revenues for the different policies depends upon the status of family 2 within a given policy. The revenue forecasts are intricate, given that some passengers in the demand forecast will prefer $f1$ over $f2$.

If the policy has both open family 1 and 2 classes, the revenue from the demand that can and will buy $f1$ is obtained, and then added to the revenue from the demand that can afford $f2$, but cannot or will not buy $f1$. If the policy only has an open family 1 class, the revenue from the demand that can afford the family 1 fare is obtained.

Recall from the passenger decision process, a passenger who can afford $f1$ chooses $f1$ only if the difference in fare between the two classes, $fare_{f1} - fare_{f2}$, is less than the disutility costs the passenger in question attributes to the family 2 restrictions. This is defined as “buy up” from family 2 to family 1. To include the probability that an eligible passenger prefers $f1$ over $f2$, or the probability of buy up from $f2$ to $f1$, the passenger's disutility distribution (assumed to be Gaussian) must be estimated. Figure 3-4 shows a random passenger's disutility distribution and how the difference in fare levels between $f1$ and $f2$ impact the likelihood of buy-up from family 2 to family 1.

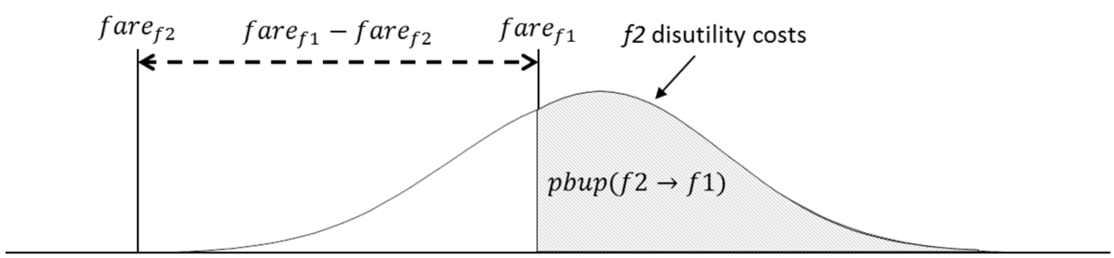

Figure 3-4: Disutility Distribution of Random Passenger. Area of Shaded Region is Probability that Passenger Buys Up from $f2$ to $f1$

The QFF1 algorithm estimates a passenger's disutility distribution through the “PBUP” parameter. The PBUP values is the fare ratio to the lowest family 2 fare at which the airline expects 50% of the eligible passengers would prefer family 1 over family 2. This input, along with an input k-factor, can then be used to generate estimated disutility distributions for all passengers. Although in practice a passenger is likely to be less sensitive to restrictions closer to departure, this is not taken into consideration in the QFF models. That is, the disutility costs for a particular passenger are assumed to be the same for each time frame throughout the booking process.

Referring back to the fare structure introduced in Table 3.1 (which has a lowest class fare of \$100), a PBUP value of 3.0 (with a disutility k-factor of 0.5) implies that the airline believes that 50% of the passengers would prefer purchasing a family 1 class priced at \$300 over a family 2 class priced at \$100. The disutility distribution can then be generated, which has a mean of \$200 (\$300 - \$100) and a standard deviation of \$100. With the estimated disutility distribution, the probability an eligible passenger will buy up from $f2$ to $f1$, denoted as $pbup(f2, f1)$, can be computed for each set of family 2/family 1 classes. Table 3.3 shows the probability of buy-up from the three family 2 classes to the three family 1 classes, assuming a PBUP input of 3.0.

From	To	$pbup(f2 \rightarrow f1)$
A2	A1	0.091
B2	A1	0.048
C2	A1	0.023
A2	B1	0.748
B2	B1	0.631
C2	B1	0.500
A2	C1	0.901
B2	C1	0.841
C2	C1	0.748

Table 3.3: Probability of Buy-Up from each Family 2 Class to each Family 1 Class

Once the disutility distributions have been estimated, the revenues can be forecasted for each of the policies. The revenue forecast equations, which depend both on the lowest open family 2 class as well as the lowest open family 1 class, are found as follows:

$$\begin{aligned} rev_{tf}\{f1, f2\} =& [FC_{tf} \times psup_{tf}(C2 \rightarrow f1) \times pbup(f2, f1)] \times fare_{f1} + \\ & [FC_{tf} \times psup_{tf}(C2 \rightarrow f2) - (FC_{tf} \times psup_{tf}(C2 \rightarrow f1) \times pbup(f2, f1))] \times fare_{f2} \end{aligned}$$

The first part of the revenue forecast equation for the policy $\{f1, f2\}$ considers those passengers who purchase the lowest open family 1 fare, that is, the travelers who both can afford and prefer $f1$. The remaining passengers who can afford $f2$ but do not choose $f1$ (either because they cannot afford or do not want $f1$) book $f2$ at a fare of $fare_{f2}$. Any passengers that do not book either $f1$ or $f2$ (i.e., passengers who can not afford $f2$) are no-go.

This approach can be used to forecast the revenues for each of the 12 policies, including the policies with only family 1 classes available. In these three policies, $pbup(f2, f1) = 1$, and the revenue forecast equation becomes

$$rev_{tf}\{f1, \%\} = FC_{tf} \times psup_{tf}(C2 \rightarrow f1) \times fare_{f1}$$

Step 4: Determine the optimal policies by constructing the convex hull

After the demands and revenues have been forecasted for each of the policies, the next step involves determining which policies are optimal. This can most easily be illustrated graphically by plotting the total revenue vs. total demand for each of the policies in a scatter plot, tracing out the convex hull (Fiig et al., 2010). Using a similar process as described in Section 2.2.3 (Fare Adjustment Theory), the optimal policies are found to be those that lie along the efficient frontier (increasing part of the convex hull), as shown in Figure 3-5. Although there are 12 policies in this numerical example, only five of them are efficient. Fiig et al. (2012) describes the main techniques used for calculating the optimal policies from the convex hull.

Step 5: Convert policy forecasts to class forecasts to obtain the optimal booking limits

Once the optimal policies for each forecast have been determined, Fiig et al. (2012) shows how, for nested policies, the policy forecasts can be converted back into class forecasts. Specifically, the adjusted fares and partitioned demand for each class can be obtained. Note that Steps 2-4 of the QFF1 process are all performed on both the Q-Forecast within the current time frame as well for the Q-Forecast for each future time frame up until departure. Thus the adjusted fares and partitioned demands for each class in each time frame are obtained.

To obtain the class bookings-to-come demand forecast, the partitioned demand

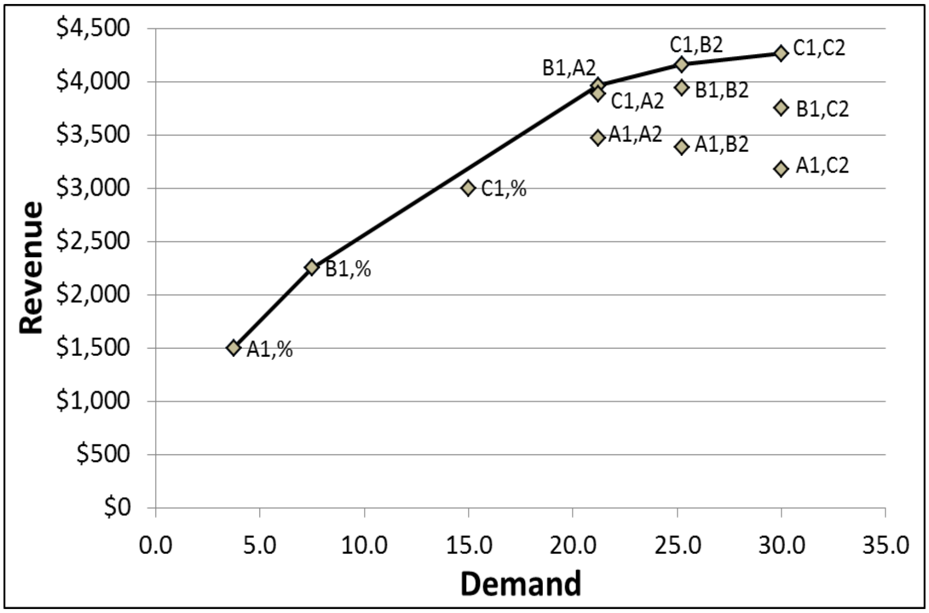

Figure 3-5: Scatter Plot of Total Revenue vs. Total Demand for all Policies

from the current time frame is summed up with the partitioned demand from the remaining time frames (for all classes). The class adjusted fare is set to the adjusted fare for the current time frame. The bookings-to-come demand forecast and adjusted fare inputs are then fed into the optimizer to produce the optimal class protection levels within the current time frame.

Limitations of QFF1

On the revenue management side, forecasting with QFF1 required both a sell-up input (FRAT5) and a buy-across input (PBUP). The single FRAT5 curve represented the airline's estimate of all passengers' WTP, while the PBUP input modeled the passengers' disutility distribution to the family 2 restrictions. While theoretically sound to include inputs that model both sell-up and buy-across, it was determined that the FRAT5 value had much larger impacts on the overall performance of QFF1 than the PBUP value. After a detailed analysis of the impacts and interactions of these two parameters on the performance of QFF1, Cizaire (2010) concluded that there was too much emphasis on the buy-up input compared to its actual role on the performance of the methodology.

Additionally, estimating both family 1 and family 2 passengers' WTP with the same sell-up model did not appear to be the most appropriate method of capturing passenger sell-up behavior. That is, with one sell-up input, there was no distinguishing between family 1 and family 2 passengers, which can have very different maximum WTP.

The reasons listed above provided the motivation to develop an alternative forecasting method for the fare family structure. This method, known as QFF2, is discussed next.

3.3.2 QFF2

A schematic illustrating the QFF2 forecasting process is shown in Figure 3-6. QFF2 puts more emphasis on forecasting separately within each family instead of forecasting for both families combined. As such, the first step includes generating separate Q-Forecasts by fare family. While the remaining steps are the same, there are notable changes in the formulations within the different steps.

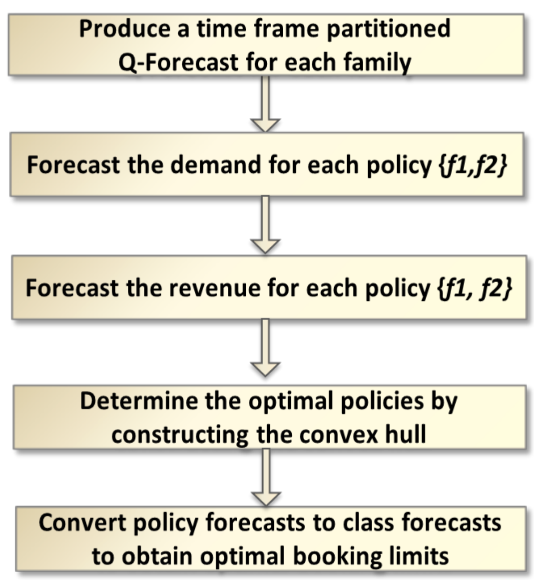

Figure 3-6: QFF2 Process

There is a major distinction between QFF1 and QFF2 in the way that sell-up is modeled. Specifically, an additional FRAT5 value is incorporated into the model. In the new formulation a separate sell-up input is used for each family, a consequence of the fact that with QFF2, separate Q-Forecasts for each family are generated in the first step.

The second difference between the methods is the removal of the buy-up input (PBUP) in the QFF2 algorithm. As mentioned earlier, it was found that the sell-up input was a dominating factor in the performance of QFF1. In the new formulation, buy-across is assumed to be implicitly taken into account, given that the forecast for each family is based on historical data available from previous bookings in each of the families. Consequently, there is no buy-up input.

In addition to the modifications mentioned above, a simplification was also made in the FRAT5 curve. Recall that in QFF1, the FRAT5 curve was modeled as an S-shape to capture the different booking tendencies between business and leisure passengers throughout the different time frames. Given that this was the only sell-up input, it was modeled in this way to account for leisure travelers (lower WTP) booking early in the process and business passengers (higher WTP) booking closer to departure.

With QFF2 producing separate forecasts by family, it is assumed that the family 1 forecast consists mainly of business travelers, while the family 2 forecast primarily consists of leisure passengers. Consequently, the separate forecasts for each family essentially segment passengers based on WTP. For forecasting purposes, it is assumed that the passengers from the family 1 forecast all have similar WTP, regardless of the different time frames. Similarly, the WTP of passengers from the family 2 forecast is not expected to change significantly as the departure date nears. Thus, QFF2 uses constant FRAT5 values by time frame instead of the S-shaped curves. While the assumptions regarding the segmentation of passengers that are listed above are not perfect, they are believed to be reasonable for forecasting purposes.

Step 1: Produce a time frame partitioned Q-Forecast for each family

As opposed to QFF1, QFF2 produces a time frame partitioned Q-Forecast separately for family 1 and family 2, both within the current time frame as well as for each future time frame before departure.

To obtain the Q-Forecast for family 1 in time frame tf ($FC1_{tf}$), only past bookings in family 1 are used. Inverse sell-up probabilities from the lowest family 1 class ($C1$) are determined from the family 1 FRAT5 value, which in this case is assumed constant throughout the booking process (e.g., 1.5).

A similar procedure is used to generate the time frame partitioned Q-Forecast for family 2 ($FC2_{tf}$). Only relevant information on past bookings in family 2 is used. Using inverse sell-up probabilities from the lowest family 2 class ($C2$), which are generated using the assumed sell-up rate (may be different from the sell-up input used in the family 1 estimate), $FC2_{tf}$ can be determined.

Step 2: Forecast the demand for each policy $\{f1, f2\}$

Recall that the FRAT5 value is the fare ratio of some higher class fare to the lowest overall fare. Since the family 1 forecast is independent of the family 2 forecast, the lowest fare for the family 1 forecast is $C1$ instead of $C2$. As a result, the sell-up equation for a passenger in the family 1 forecast ($psup1$) becomes

$$psup1_{tf}(\text{C1} \rightarrow f) = e^{(\frac{fare_f}{fare_{C1}}-1)(\frac{ln(0.5)}{FRAT5_{tf}-1})}$$

Given that the lowest family 1 fare is always priced higher than the lowest family 2 fare, the same numerical sell-up input (e.g., 2.0) will implicitly produce steeper sell-up estimates in family 1 than in family 2. For example, in the fare structure from Table 3.1, a FRAT5 value of 2.0 in both families estimates a median WTP of \$400 for family 1 passengers and a median WTP of \$200 for passengers from the family 2 forecast.

The demand forecasts with QFF2 for each policy are straightforward, given that the forecasts for each family are independent from one another. The maximum demand forecast for any policy $\{f1, f2\}$ is found as follows:

$$dem_{tf}\{f1, f2\} = FC1_{tf} \times psup1_{tf}(C1 \rightarrow f1) + FC2_{tf} \times psup2_{tf}(C2 \rightarrow f2)$$

In this formulation, the maximum demand for any policy is found by estimating the number of passengers from the family 1 forecast that can afford $f1$, and then adding this number to the estimated number of passengers from the family 2 forecast that can afford $f2$. This model assumes that there is no interaction between the two families. That is, no passengers from the family 2 forecast sell-up to $f1$ in the algorithm. Likewise, any passengers from the family 1 forecast who do not sell up to $f1$ are assumed to be no-go (rather than book $f2$).

Step 3: Forecast the revenue for each policy $\{f1, f2\}$

The revenue forecasts for each policy in the QFF2 formulation are also calculated in a straightforward manner, given that buy-across is not captured in the model, and thus there is no probability of buy-up from family 2 to family 1 in the formulation. As a result, the algorithm assumes that only passengers in the family 1 forecast buy $f1$, and only passengers in the family 2 forecast buy $f2$. Specifically, the revenue forecast for any policy $\{f1, f2\}$ is

$$\begin{aligned} rev_{tf}\{f1, f2\} = &[FC1_{tf} \times psup1_{tf}(C1 \rightarrow f1)] \times fare_{f1} + \\ &[FC2_{tf} \times psup2_{tf}(C2 \rightarrow f2)] \times fare_{f2} \end{aligned}$$

That is, the revenue forecast is calculated using the demand for $f1$ and $f2$, multiplying these values by $fare_{f1}$ and $fare_{f2}$, respectively, and then adding the results. This process can be done for each of the possible policies an airline can offer.

Steps 4-5

The remaining steps are identical to the QFF1 algorithm. All the policies are plotted in a scatter plot of total revenue vs. total demand. Assuming the policies on the efficient frontier are nested, the policy forecasts can be mapped back to class forecasts to eventually obtain the booking limits.

Limitations of QFF2

QFF2 differs from the original fare family forecasting method in two ways. Separating the forecasts by fare family and treating each family independently appeared to be more accurate from a theoretical standpoint. Under this assumption, buy-across was not modeled, and the buy-up input was removed. However, several experiments using the PODS simulator showed that, in practice, a certain percentage of passengers from the family 2 forecast bought a family 1 product, and vice versa. The QFF2 algorithm did not capture this interaction between the two families.

The third QFF formulation (QFF3) was developed with the intention of incorporating the best characteristics of the first two models into a new formulation.

3.3.3 QFF3

The two forecasting methods discussed thus far utilized different approaches in modeling the passenger decision process in a fare family environment. The QFF1 formulation captured both sell-up and buy-across. One of the drawbacks to this method was estimating the sell-up probabilities of both passenger types with a single sell-up model. This provided part of the motivation for QFF2, where the passengers for each family are forecasted separately, including different sell-up inputs by family. Buy-across was not included in the model, and as a result, there was no buy-up input. However, experiments showed that there was still interaction between the families that was not being modeled in the QFF2 algorithm. The third formulation is essentially a hybrid between QFF1 and QFF2.

QFF3 includes separate forecasts and requires separate sell-up estimates by family (as in QFF2). As such, the only differences between QFF2 and QFF3 are in the formulations. With QFF3, buy-across is modeled in the form of buy-up as well as "buy-down", given the separate family forecasts. With QFF3, the probability of buy-up is the probability a passenger from the family 2 forecast prefers $f1$ over $f2$. On the flip side, the probability of buy-down is the probability that a passenger in the family 1 forecast prefers $f2$ over $f1$.

With QFF3, an airline can estimate the disutilities attributed to the family 2 restrictions through the use of the new disutility multiplier input "DUMLT". The mean estimated disutility costs for a passenger in the family 1 forecast ($DUMU_1$) and a passenger in the family 2 forecast ($DUMU_2$) are

$$DUMU_1 = DUMLT \times fare_{C1}$$

$$DUMU_2 = DUMLT \times fare_{C2}$$

In this formulation, a single DUMLT value is used to estimate the disutilies attributed to the family 2 restrictions for both family 1 and family 2 passengers. This is accomplished by multiplying the DUMLT input times the price point of the lowest booking class in each family. Finally, the disutility k-factor (assumed to be 0.5) can be used to define the standard deviation of the disutility distributions, which are again assumed to be Gaussian.

With these estimated disutility distributions for both passengers from the family 1 forecast and the family 2 forecast, the probability of a passenger choosing a booking class in the opposite family in which he was forecasted can be determined. To this end, the probability a passenger in the family 1 forecast would buy down to the lowest open family 2 fare class $f2$ is

$$pbdwn(f1, f2) = P(DUMU_1 < f1 - f2)$$

That is, the probability that a passenger in the family 1 forecast would prefer the lowest family 2 class is the probability that the difference in fare between $f1$ and $f2$ is greater than the estimated disutilities as defined above.

In a similar manner, the probability of a passenger in the family 2 forecast buying up to $f1$ can be found as follows:

$$pbup(f2, f1) = P(DUMU_2 > f1 - f2)$$

The additional parameters added were done so to explicitly represent buy-up and buy-down, and thus better model buy-across in the actual passenger decision process. These parameters are incorporated into the QFF3 revenue forecast. Note that the only changes between the QFF2 and QFF3 formulation lies in the demand and revenue forecasts (Steps 2-3). As such, these are the only steps from the QFF forecasting process that are highlighted in the description of the QFF3 algorithm.

Step 2: Forecast the demand for each policy $\{f1, f2\}$

The demand forecast with QFF3 is dependent upon the status of the lowest open family 2 class. Assuming some family 2 class is open, the demand forecast for any policy $\{f1, f2\}$ is

$$dem_{tf}\{f1, f2\} = FC1_{tf} \times psup1_{tf}(C1 \rightarrow f2) + FC2_{tf} \times psup2_{tf}(C2 \rightarrow f2)$$

In this equation, $psup1_{tf}(C1 \rightarrow f2)$ is the probability a family 1 passenger will sell-up to $f2$ in time frame tf, while $psup2_{tf}(C2 \rightarrow f2)$ is the probability a passenger from the family 2 forecast will sell-up to $f2$. The main difference between the two quantities is that the sell-up probability for the passenger from the family 1 forecast is computed with $C1$ as the lowest class, while the latter is computed with $C2$ as the lowest class.

The total demand for any policy with a family 2 class open is the number of passengers from both the family 1 forecast as well as the family 2 forecast that can afford $f2$. In a non-overlapping fare family structure, every passenger in the family 1 forecast is included in the demand, given that each passenger is assumed to be able to afford $f2$.

If family 2 is completely closed, the demand forecast with QFF3 becomes

$$dem_{tf}\{f1, \%\} = FC1_{tf} \times psup1_{tf}(C1 \rightarrow f1) + FC2_{tf} \times psup2_{tf}(C2 \rightarrow f1)$$

That is, the forecast is now the number of passengers from both the family 1 forecast and family 2 forecast that can afford $f1$.

Step 3: Forecast the revenue for each policy $\{f1, f2\}$

With QFF3, buy-across includes both passengers from the family 2 forecast buying up to family 1 as well as passengers from the family 1 forecast buying down to family 2. These passenger characteristics, along with sell-up, are all included in the QFF3 revenue forecasts. To this end,

$$rev_{tf}\{f1, f2\} = rev(f1) + rev(f2)$$

The forecasts for the lowest open family 1 class as well as the lowest open family 2 class depend both on $FC1_{tf}$ and $FC2_{tf}$, that is, the Q-Forecasts from each family in time frame tf. The revenue forecast for the lowest open family 1 class is

$$\begin{aligned} rev_{tf}(f1) =& [FC1_{tf} \times psup1_{tf}(C1 \rightarrow f1) \times (1 - pbdwn(f1, f2))] \times fare_{f1} + \\ & [FC2_{tf} \times psup2_{tf}(C2 \rightarrow f1) \times pbup(f1, f2)] \times fare_{f1} \end{aligned}$$

The first part of the equation consists of the passengers from the family 1 forecast that purchase $f1$. This includes those passengers from the family 1 forecast who can both afford $f1$ and prefer $f1$ (i.e., those that do not buy down to family 2). The second component of the equation consists of those passengers from the family 2 forecast who can both afford and prefer $f1$. The number of passengers forecasted to purchase $f1$ (from both the family 1 and family 2 forecasts) is multiplied by $fare_{f1}$ to obtain the forecasted revenue contribution to the policy $\{f1, f2\}$ from $f1$.

Using similar logic, the revenue forecast for the lowest open family 2 class can be calculated as

$$\begin{aligned} rev_{tf}(f2) =& [FC1_{tf} \times psup1_{tf}(C1 \rightarrow f2) - FC1_{tf} \times psup1_{tf}(C1 \rightarrow f1) \times (1 - pbdwn(f1, f2))] \times fare_{f2} + \\ & [FC2_{tf} \times psup2_{tf}(C2 \rightarrow f2) - FC2_{tf} \times psup2_{tf}(C2 \rightarrow f1) \times pbup(f2, f1)] \times fare_{f2} \end{aligned}$$

The revenue forecast for $f2$ considers both those passengers who buy-down from the family 1 forecast as well as passengers from the family 2 forecast. The majority of passengers from the family 1 forecast are assumed to purchase $f1$. Those who don't purchase in $f1$ either purchase $f2$ or are no-go. This is captured in the above formulation by $FC1_{tf} \times psup1_{tf}(C1 \rightarrow f2)$. Note that $psup1_{tf}(C1 \rightarrow f2) = 1$ in a non-overlapping fare structure, since in this case all passengers in the family 1 forecast are expected to be able to afford any family 2 class, given that the lowest family 1 class is priced above the most expensive family 2 class. Note that, in an overlapping fare family structure, the probability a passenger from the family 1 forecast sells up to one of the more expensive family 2 classes can be less than 1, given that some family 2 classes are priced above the lowest family 1 class.

The final part of the revenue equation for $f2$ comes from the passengers from the family 2 forecast. This includes all passengers from $FC2_{tf}$ that can afford $f2$ but do not buy $f1$. Finally, the values obtained are multiplied by $fare_{f2}$ to produce the revenue contribution to the policy from $f2$.

QFF3 summary

It is believed that the QFF3 formulation appropriately captures the different aspects of the passenger decision process in a fare family structure. The first two formulations had limitations, which led to the development of subsequent methods. From a formulation perspective, there are no apparent invalid assumptions in the model. It is the three methods described throughout this chapter that are tested and analyzed in this thesis.

3.4 Chapter Summary

This chapter provided a detailed description of the fare family forecasting algorithms. The assumed passenger decision process in a fare family structure was described, providing the basis for the developed methods. Other essential terms and concepts were introduced in more detail, including passenger sell-up and buy-across. An in-depth walk-through of a numerical example with QFF1 was illustrated, followed by an analysis of the differences between the original fare family forecasting method and each of the subsequent formulations.

Following the descriptions of each of the first two methods, the limitations of each were discussed. With QFF1, it was found that, although buy-across was accounted for, it had little impact on overall performance. Additionally, there was motivation to split up the forecasts by family, which led to the development of QFF2. In the second formulation, there were two sell-up inputs and no buy-across input. While producing separate forecasts appeared to more accurate from a modeling perspective, passengers from one family forecast were still purchasing in the opposite family, which was not captured in the algorithm.

The last method, QFF3, was designed with the goal of incorporating the best characteristics from each of the first two methods into one formulation. Specifically, separate inputs to estimate sell-up were still required for each family in the third formulation. However, unlike with QFF2, buy-across was accounted for with QFF3 through the DUMLT parameter. This input was used to estimate buy-across in the form of buy-up and buy-down.

The next chapter introduces the RM simulator used to test the different experiments throughout this thesis, the Passenger Origin-Destination Simulator. The airline networks used to set up the experiments and test the performance of the different algorithms are described in detail.

Chapter 4

Overview of the Passenger Origin Destination Simulator (PODS)

This chapter introduces the Passenger Origin Destination Simulator (PODS), the software tool used to test and analyze the performance of the different RM methods that were described in the previous chapter. PODS was developed in the mid-1990s as a successor to The Boeing Company's Decision Window Model, which originally was created to model passenger choice with respect to flight schedules. The simulator has been further advanced by the MIT PODS Consortium, funded by several members associated with the airline industry. PODS simulates hypothetical airline networks, in which different RM system forecasting and optimization methods' performances can be analyzed. The first section in this chapter briefly describes the main components of the PODS architecture. This overview includes details of the Passenger Choice Model, the Revenue Management System, and the interaction between these two components. The simulated two-airline environments used for analysis will be introduced in the second section. For more information on the PODS simulator, see Belobaba (2010)

4.1 PODS Architecture

In PODS, each simulation run begins with user defined inputs. The initial iterations allow the airlines to progressively build the historical database, which is used to generate forecasts. A typical PODS run is composed of five independent trials, with each trial consisting of 600 samples, or "departure days". The results from a single run are obtained by discarding the first 200 samples of each trial to eliminate the initial condition effects, and then averaging over the 400 samples from each of the five trials (total of 2000 iterations). The results are based on simulated individual passenger choices and convey information on various airline performance metrics, including airline traffic, revenues, load factors, and yields.

The pre-departure process in PODS is modeled as a 63-day process, where passengers can begin booking up to 63 days prior to departure. These days are divided into 16 time frames, where the RM system updates seat availabilities by class at the beginning of each time frame. As Table 4.1 illustrates, the length of each time frame is initially about seven days due to fewer anticipated bookings in the early stages, but then is compressed to only a few days closer to departure.

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

Table 4.1: Booking Period Time Frames

The passenger choice process in PODS is captured in the Passenger Choice Model (enhanced version of Decision Window Model), where passengers choose between different airlines, booking classes, and itineraries. The different path class availabilities generated by an airline's Revenue Management System are fed into the passenger choice model. Figure 4-1 illustrates how, although the main components of the PODS architecture are separate, both interact with one another. These two pieces of the PODS puzzle are now discussed in greater detail.

4.1.1 Passenger Choice Model

The Passenger Choice Model in PODS models passenger behavior through four separate steps: (1) Demand Generation; (2) Passenger Characteristics; (3) Passenger Choice Set; and (4) Passenger Decision. Each of these steps will briefly be introduced.

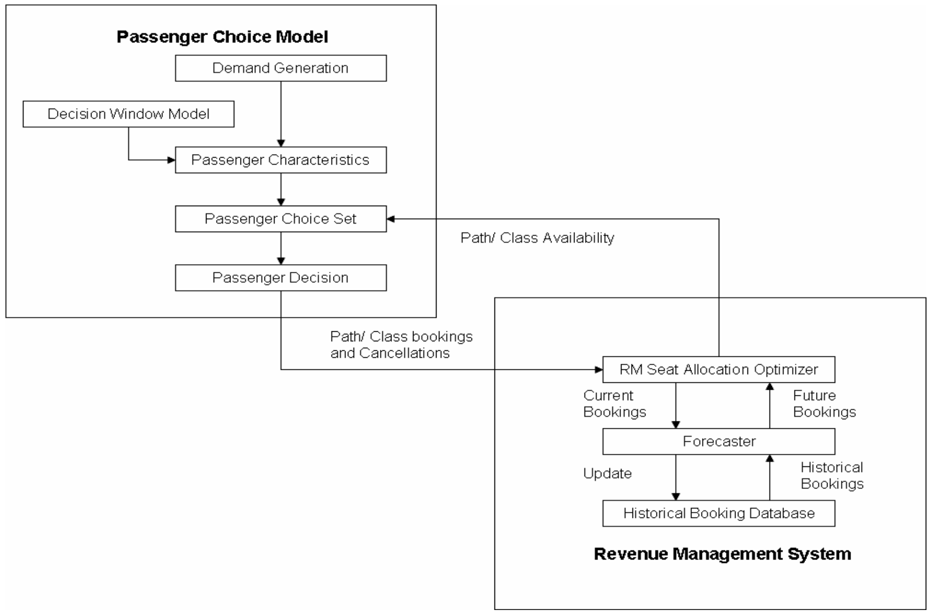

Figure 4-1: PODS Architecture (Belobaba, 2010)

Please refer to Carrier (2003) for a more in-depth overview of the model.

4.1.1.1 Demand Generation

For any given PODS scenario, the average total demand for each of the OD markets in the network is input by passenger type. Passengers are classified as either ‘business’ or ‘leisure’, with business passengers comprising about 35% of the total demand (percentages based on input from members of the PODS Consortium). The Passenger Choice Model then generates variability around the average demand for the different OD markets to draw the actual demand on a given day. While variability is added randomly to the demand generated, seasonal variability and other trends are typically not accounted for. Once the demand for each market is established for a single sample, the arrival of business and leisure travelers is set according to user-defined booking curves. The booking curves used in this thesis are shown in Figure 4-2. As the figure illustrates, past history shows that less time-sensitive leisure passengers tend to arrive much earlier in the booking process than than business travelers.

In PODS, for each passenger type there are two inputs which impact the demand generation process: the “base fare” and the “demand multiplier”. For any OD market,

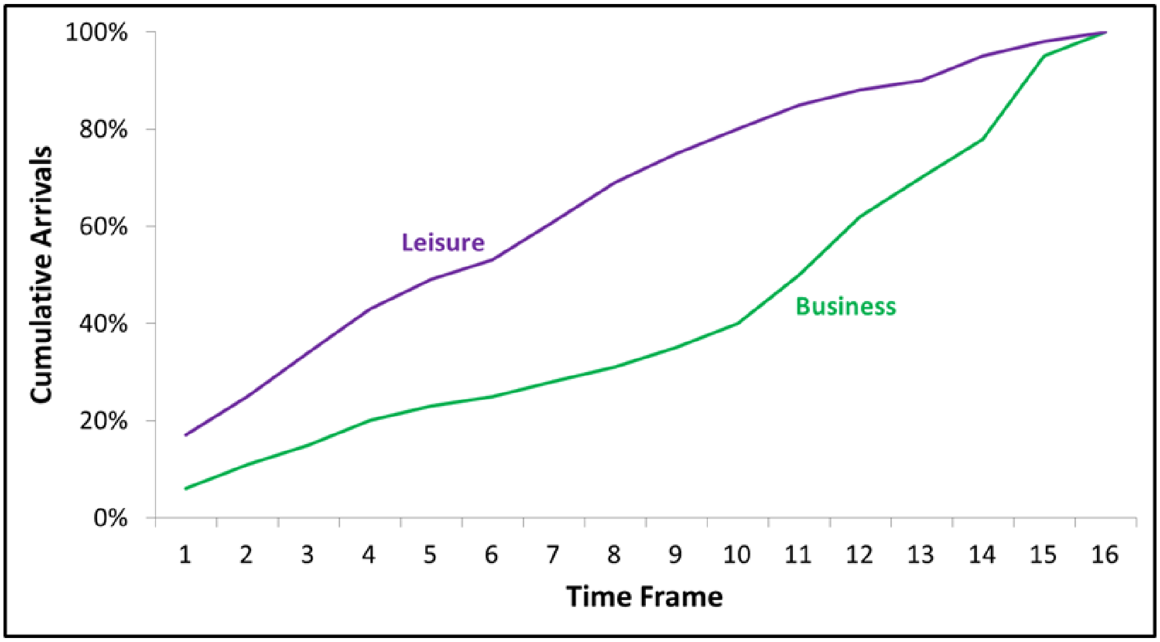

Figure 4-2: Booking Curves in PODS by Passenger Type

the base fare for that market is defined as the price point at which the mean number of passengers input for that market would be willing to fly. By adjusting the base fare, the user has the ability to increase or decrease the demand for that market. Another way to alter the demand is by modifying the demand multiplier, a setting in PODS used to generate higher or lower demand by season. In this thesis, the demand multiplier was set to generate baseline cases with approximately 85% load factor, although higher and lower demand levels can also tested by adjusting the demand multiplier.

4.1.1.2 Passenger Characteristics

Three sets of characteristics are assigned to passengers in the Passenger Characteristics step: a decision window, a maximum WTP, and a set of disutility costs. The decision window is different for each passenger, and is defined as the period from the earliest acceptable departure time to the latest allowable arrival time. Given that leisure travelers are less time-sensitive, they are assigned a wider decision window than the business passengers. Any time outside a passenger's decision window will contribute to the associated product's disutility costs.

PODS generates a maximum WTP for each passenger, that is, a maximum fare a given passenger is willing to pay to travel. This WTP is generated from user-defined inputs using the following formula:

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

In the above equation, f is the fare of the travel alternative, $basefare$ is the input in PODS representing the price which the mean number of leisure passengers input for each OD market are willing to pay to travel, and $emult$ is the elasticity multiplier, such that 50% of passengers are willing to pay $emult \times basefare$ to travel. The WTP probabilities for each passenger type are shown in Figure 4-3. Note that the base fare for business passengers in each market is 2.5 times the base fare for leisure travelers. Thus the user can control both leisure and business passengers' WTP by modifying the $basefare$ input for each market.

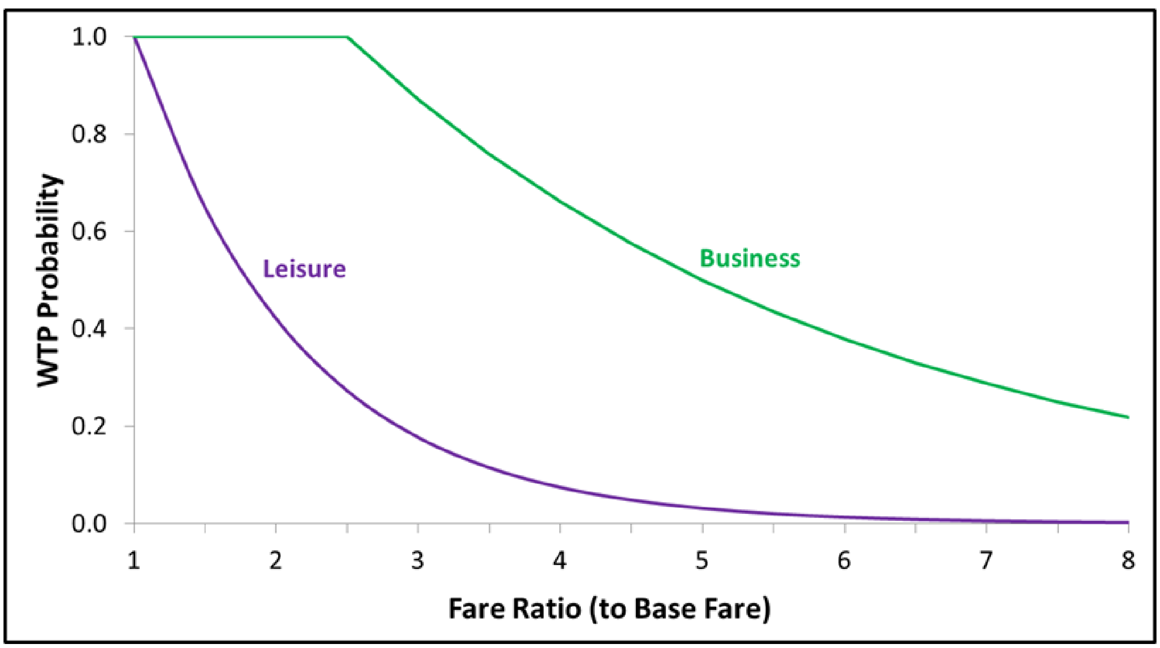

Figure 4-3: WTP Curve by Passenger Type

Lastly, disutility costs are randomly generated and assigned to each passenger type. In PODS, there are capabilities to represent different types of disutilities, including passengers' sensitivity to schedule preference (if original itinerary was outside decision window), path quality (non-stop vs. connecting), airline preference, Saturday night minimum stay requirements, itinerary change fees, and non-refundability constraints. For more information on modeling passenger disutilities, consult Lee (2000).

4.1.1.3 Passenger Choice Set

Once a given passenger is assigned a full set of characteristics, he is then presented with a set of different fare product options, along with the no-go option as well. In PODS, whenever there is some fare product that meets all of his criteria, he will not select the no-go option. Of all the options that were originally included in the passenger's decision window, some of them will be removed due to the class no longer being available (due to RM closures), an advanced purchase requirement that cannot be met, or a price point above the passenger's maximum WTP.

4.1.1.4 Passenger Decision

The passenger is then given a choice among the available alternatives. Each of the options are ranked according to their total generalized cost (as described in Chapter 3.2). The passenger will choose the option that has the lowest total generalized cost. Once the decision is made, the booking is recorded into the airline RM system's historical booking database.

4.1.2 Revenue Management System

Thus far, only the Passenger Choice Model of the PODS architecture has been discussed. The airline component of PODS consists of a third-generation RM system, similar to the RM system described in Figure 2-1. That is, the RM system is made up of a Historical Bookings Database, a Forecaster, and a Seat Allocation Optimizer.

The historical bookings database records the path/class of every booking on each airline in the network. The default bookings used for the beginning of the simulation are replaced by actual observed bookings as the simulation progresses.

The information from the historical bookings database is used by the forecaster to provide a forecast of future demand by leg or path. Because the historical demand only includes passengers who actually flew (and not necessarily those that were turned away due to no availability), the booking data is unconstrained prior to forecasting the future demand.

4.1.2.1 Seat Allocation Models in PODS

Although PODS has a number of different seat allocation optimizers, representative versions of EMSRb leg control and DAVN network RM are the only ones used throughout this thesis. These seat allocation optimizer takes the forecasts generated as inputs and determines the availability of booking classes on every leg/path. It is worth noting that the main purpose of this thesis is not so much to compare the accuracy of EMSRb and DAVN, but rather to compare the performance of the various forecasting optimization methods discussed in the previous section.

4.1.2.2 Forecasting Methods in PODS

In Section 2.1, some of the more common forecasting methods used in practice were introduced. In this section, the forecasting methods which will be utilized in this thesis are briefly discussed. Pick-up Forecasting (as described in Chapter 2) is used as the standard forecasting method throughout this thesis, and is referred to as "Standard Forecasting" for the remaining chapters. Note that with Standard Forecasting, the demand for each time frame can be forecasted at either a leg/class level or path/class level. In this research, these methods are referred to as Standard Leg Forecasting and Standard Path Forecasting, respectively.

Both Q-Forecasting and Hybrid Forecasting, developed under the research of the PODS consortium, were also described in Chapter 2. These methods utilized the notion of sell-up. In PODS, an airline's estimate of the sell-up probabilities is governed by the FRAT5 input, which was defined as the fare ratio of a higher fare to the lowest fare class at which 50% of the demand for the lowest booking class would be willing to sell-up to the higher class. That is, this parameter is an estimate of what the airline expects from passengers' regarding their sell-up behavior (Belobaba, 2010). Less price-sensitive passengers are generally expected to have a higher FRAT5 than those passengers that are more price-sensitive (i.e., business travelers).

In PODS simulations, the typical FRAT5 curve is known as "FRAT5c", which has values ranging from 1.2 in the first time frame up to 3.0 in the final time frame. Other sell-up inputs in PODS include the more-aggressive "FRAT5a" as well as the less-aggressive "FRAT5e". These FRAT5 curves are all plotted in Figure 4-4. As the figure illustrates, the distinct S-shaped curve is much more visible in the more-aggressive FRAT5 curves. It is worth mentioning that these sell-up curves are only used with QFF1 in this thesis, as constant sell-up inputs are assumed for both QFF2 and QFF3 (as mentioned in the previous chapter).

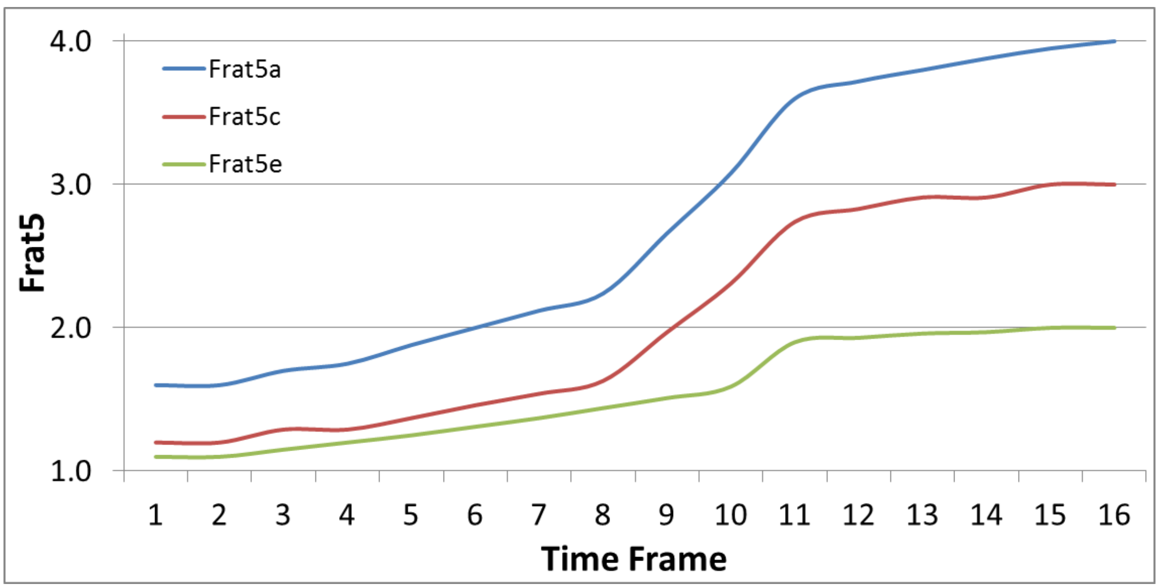

Figure 4-4: Different FRAT5 Curves used in PODS

4.2 Simulation Environment

There are various competitive airline networks in PODS that can be used as the basic structure for setting up a simulation experiment. Because different revenue management methods are often designed to operate in a variety of environments, it is useful to test these methods under multiple conditions. The conditions are accounted for by the numerous different controllable attributes for each network in the simulation. The most important factors to consider in the experiments are the airlines' fare structures with various restrictions, as well as the size and complexity of the network and markets served by the airline of interest. This thesis makes use of a single network structure in PODS known as Network D10.

4.2.1 Network D10

Network D10 in PODS is a dual airline competitive network, with each airline serving 40 spoke cities out of its central hub. All traffic flows from west to east, and connects through one of the two hubs. A total of 482 distinct OD markets are served by the two airlines; each airline operates 126 legs per day. A map of the network is shown in Figure 4-5. All of Airline 1's (AL1) legs either arrive at or depart from the AL1 hub (H1), while all of Airline 2's (AL2) legs arrive at or depart from H2. Throughout the PODS simulations in this thesis, the experiments are set up so that AL1 is the airline of interest, leaving AL2 as the competitor.

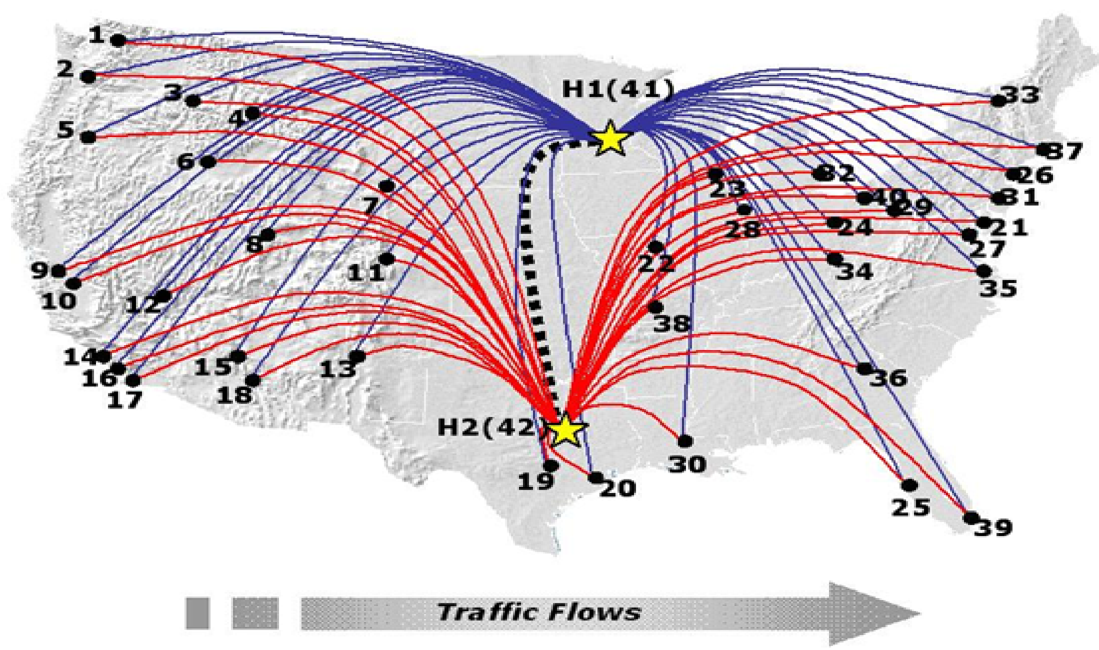

Figure 4-5: Map of Network D10

PODS offers the capability to control which restrictions an airline imposes for the different booking classes. These restrictions include the Saturday night minimum stay requirement, cancellation fee, and the non-refundability of a ticket. In this research, both airlines are assumed to offer identical fare structures, that is, two identical fare families. The families are differentiated from one another by both service and flexibility. As shown in the previous chapter, within each family in a fare family structure the booking classes are undifferentiated with respect to the imposed restrictions. The only difference between classes in each family is the price points, which are in decreasing fare order. The relatively symmetric network was chosen to ensure that the main distinction between the two competing airlines is their RM system, thus allowing a better evaluation of the performance of the different forecasting and optimization methods discussed in the previous chapter.

Within a fare family setup there are two different types of fare family structures: non-overlapping and overlapping structures. In a non-overlapping fare family structure, the least expensive family 1 class is priced above the most expensive family 2 class. The fare structure used in the numerical example in the previous chapter (Table 3.1) is an example of a non-overlapping fare structure. As its title suggests, an overlapping fare family structure exists when there is some overlap in price points between the two families. That is, at least one of the lower family 1 classes is priced below at least one of the top family 2 classes. The two fare family structures used in this thesis are introduced in the following sections.

4.2.1.1 Non-Overlapping Fare Family Structure

The booking classes for the fare family structures are labeled the same way as in the numerical example from chapter 3, with A1 being the most expensive family 1 class and A2 being the highest-priced family 2 class. However, unlike the example which included only three classes in each family, the fare structure used in the experiments in this research include five classes in each family. As such, the cheapest classes in each family are E1 and E2 for family 1 and family 2, respectively. These booking classes, along with their associated price points and restrictions, are shown in Figure 4-6. The fares under the “average fare” column are averaged over all 482 markets. In the network used throughout this research (i.e., Network D10), the fares for each market are generally proportional to the distance of each individual OD market, but vary from market to market randomly.

Fare Family	Class Name	Restrictions: Sat. Stay	Restrictions: Cancel Fee	Restrictions: No Refund	Avg. Fare ($)
Unrestricted – Family 1	A1	✗	✗	✗	500
	B1	✗	✗	✗	450
	C1	✗	✗	✗	400
	D1	✗	✗	✗	350
	E1	✗	✗	✗	300
Restricted – Family 2	A2	✓	✓	✓	220
	B2	✓	✓	✓	190
	C2	✓	✓	✓	160
	D2	✓	✓	✓	131
	E2	✓	✓	✓	101

Figure 4-6: Fare Family Structure with Non-Overlapping Price Points

4.2.1.2 Overlapping Fare Family Structure

In addition to testing the different forecasting and optimization methods with the fare family structure shown in Figure 4-6, experiments were also conducted on an alternative fare structure, where one or more of the classes from family 2 are priced higher than the cheapest family 1 class.

To develop the overlapping fare family structure, the non-overlapping fare family structure introduced in the previous section was modified. The fare levels of the family 1 booking classes were lowered while the family 2 prices were increased. The overlapping fare structure is shown in Figure 4-7. From this structure arise two important insights that should be noted. First, the price points of the lowest two booking classes did not change, and thus no new demand was stimulated from before.

Second, the price points of the top two family 2 classes are now more expensive than class E1, creating the intertwined fare structure.

Fare Family	Class Name	Sat. Stay (Restrictions)	Cancel Fee (Restrictions)	No Refund (Restrictions)	Avg. Fare ($)
Unrestricted – Family 1	A1	✗	✗	✗	500
	B1	✗	✗	✗	425
	C1	✗	✗	✗	350
	D1	✗	✗	✗	275
	E1	✗	✗	✗	200
Restricted – Family 2	A2	✓	✓	✓	280
	B2	✓	✓	✓	220
	C2	✓	✓	✓	170
	D2	✓	✓	✓	131
	E2	✓	✓	✓	101

Figure 4-7: Fare Family Structure with Overlapping Fares in Network D10

4.3 Chapter Summary

This chapter introduced the simulation environment used to test and analyze the performance of the fare family forecasting methods described in chapter 3. The basic architecture of PODS was described, including the Passenger Choice Model and the Revenue Management System. The Passenger Choice Model provides the user with tremendous flexibility in designing experiments to test different RM methods. The Revenue Management System's chosen for this thesis were modeled as a typical third-generation RM system.

The airline network (Network D10) used for all of the simulations was described in detail. This two airline network allows us to effectively analyze the performance of one airline while using the opposite airline as the competitor.

Finally, the two-family ten-class fare family structures used in this thesis were introduced. The non-overlapping fare family structure is used in the majority of the simulations throughout this thesis. However, the performance of the QFF methods is also tested with in overlapping fare structure to determine how relative performance is impacted.

In the next chapter, the results from the PODS simulations that tested the different forecasting and optimization methods from chapter 3 are presented.

Chapter 5

PODS Simulation Results

This chapter presents the results from the different experiments that were performed to test the forecasting and optimization methods developed for fare family structures. The QFF methods described in Chapter 3 are analyzed in detail using different metrics (e.g., revenues, load factors) to better understand the strengths and drawbacks of each model. All simulations in this chapter are performed in the dual-airline Network D10 (as described in Chapter 4).

This chapter is divided into two sections. First, each QFF method is tested with a fare family structure with non-overlapping price points (denoted as "non-overlapping structure"), as previously shown in Figure 4-6. Both leg-based and origin-destination (OD) RM controls are tested with the different fare family forecasting methods. The second section illustrates the performances of the different QFF methods when overlapping fare family structures ("Overlapping") are implemented for both airlines (see Figure 4-7 for the average fares for each class in this structure).

Each experiment performed in this chapter is categorized as either a "symmetric RM" or a "competitive RM" test case. Although the two airlines in the network are competitors, the symmetric RM environment is defined as the experimental setup in which both airlines in the network are assumed to use identical seat allocation models, forecasting methods (including any sell-up/buy-across inputs), and advanced purchase requirements (or lack thereof). Although this type of scenario is unlikely in practice, testing forecasting and optimization methods in the symmetric environment within PODS provides the ability to better analyze the behavior of the algorithm without any competitive feedback effects. That is, this type of setup is in many respects the same as a single airline environment.

In all competitive RM experiments, the seat allocation model and/or forecasting method for Airline 1 is modified while holding the competitor's (Airline 2) RM methods constant. In these experiments, basic RM controls for the competitor airline are assumed in all competitive RM simulations; Airline 2 uses a leg-based seat allocation optimizer (EMSRb), Standard Forecasting, and advanced purchase (AP) requirements on the different classes. These AP restrictions are shown in Table 5.1. As the values in the table indicate, a slightly more aggressive AP is used for family 2, with the lowest E2 class closing down (at the latest) three weeks prior to departure. The top classes in each family have no AP restriction.

Class	AP
A1	0
B1	3
C1	7
D1	10
E1	14
A2	0
B2	7
C2	10
D2	14
E2	21

Table 5.1: Advanced Purchase Requirements used by Airline 2 in Competitive RM Simulations

5.1 Non-Overlapping Fare Family Structures

The first section of this chapter provides an analysis on the results from all the simulations in which both airlines used fare family structures with non-overlapping price points. The subsections contain the results from the simulations in which both leg-based and OD controls are used in the symmetric RM environment. Following this, the results from the competitive RM experiments are presented. In each section, the performances of the different forecasting and optimization methods tested throughout this thesis are analyzed relative to a "base case". Different base cases will be used depending on the simulation environment (i.e., competitive RM vs. symmetric RM) as well as the types of controls used (leg-based vs. OD). In each section, the appropriate experimental setup for the base case will be specified, along with the corresponding results from the experiment in question.

5.1.1 Leg-based Controls in Symmetric RM Experiments

The base case used throughout this section is the experimental setup in which both airlines implement EMSRb with Standard Forecasting and an advanced purchase fare family structure. The AP requirements are identical for both airlines, and is the same set of restrictions that Airline 2 uses in the competitive RM simulations (Table 5.1). Although no AP is used in conjunction with QFF (as described in Chapter 3), it is used with Standard Forecasting in the base cases to prevent the airlines' RM systems from spiraling down. In all leg-based symmetric RM test cases, Airline 2 implements the same RM methods as Airline 1.

In addition to Standard Forecasting, Hybrid Forecasting with FRAT5c (and an AP requirement) was also tested in the preliminary leg-based simulations. The results from the symmetric RM experiments in which both airlines use Standard Forecasting (i.e., the base case), as well as Hybrid Forecasting (HF), are shown in Figure 5-1. In the base case, Airline 1 achieves a total revenue (over the entire network) of $1,388,833. The revenues obtained from subsequent leg-based symmetric RM simulations are all measured relative to this amount. The baseline load factor is 86% for both airlines.

As expected, Hybrid Forecasting provides an increase in revenue over Standard Forecasting in a symmetric RM environment. Specifically, Hybrid Forecasting leads to a 4.5% increase in revenue for Airline 1 over the base case. The load factor for each airline decreases to 81% with Hybrid Forecasting, owing to the fact that the higher revenues have been obtained through more effective lower-class closures,

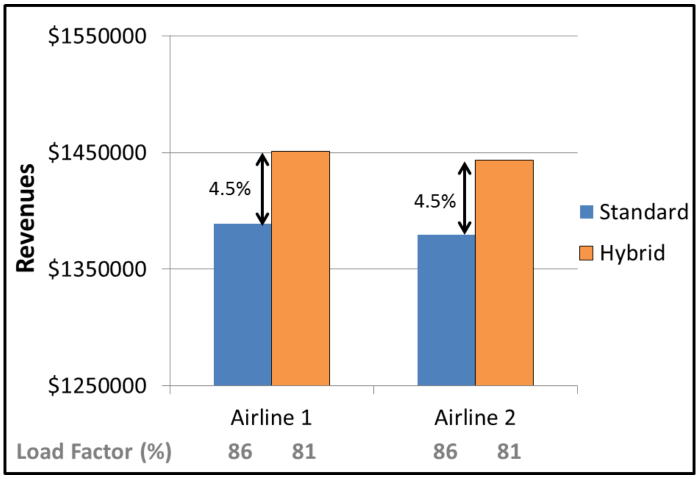

Figure 5-1: Baseline Revenues

causing passengers to sell up to higher-priced classes. Despite this gain in revenue over the base case, there is little scientific evidence that Hybrid Forecasting appropriately models the passenger decision process in a fare family structure. It is worth noting that, given the RM symmetry between the two airlines, it is not surprising to observe that both airlines perform very similar to one another in the different metrics shown.

The number of passengers booked in each class over the entire network, that is, the "booking class mix", is another metric used frequently throughout this chapter. The number of passengers in a particular class can be further broken down into the different types of passengers (i.e., business or leisure) to better understand the different passenger behaviors.

To illustrate the mix of passengers in the baseline experiments, the total number of business and leisure bookings in each class over all OD markets for Airline 1 is shown in Figure 5-2(a) for Standard Forecasting. As expected, the majority of family 1 passenger bookings are from business travelers (91%). Overall, business passengers represent approximately 38% of the total bookings. Due to the effects of the AP requirements, bookings occur in all 10 classes, with a fairly significant amount of business travelers buying down to the highest family 2 class (A2).

The business/leisure mix with Hybrid Forecasting (for Airline 1) is given in Figure 5-2(b). The increase in revenue obtained with Hybrid Forecasting is a direct result of the higher quantity of bookings in each of the top top eight classes. That is,

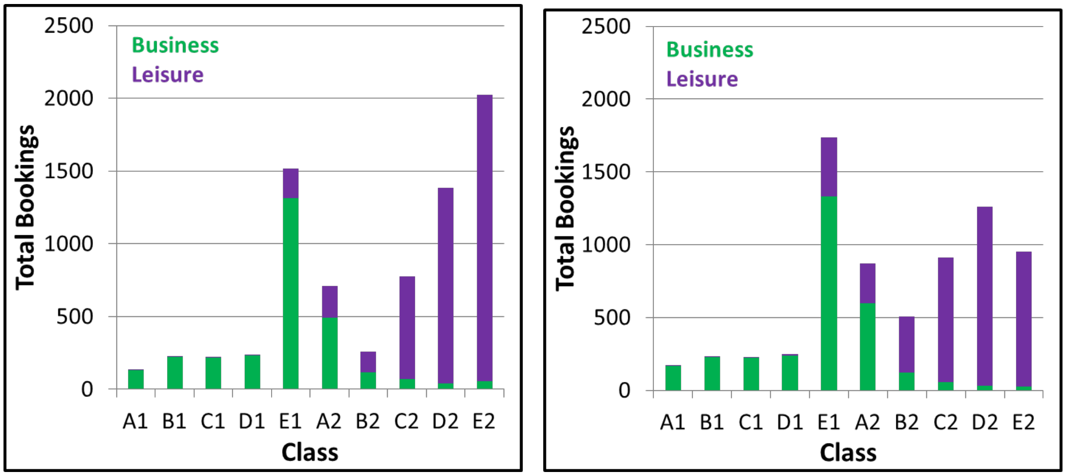

(a) Standard Forecasting

(b) Hybrid Forecasting

Figure 5-2: Baseline Business/Leisure Bookings for Airline 1

closing E2 more effectively and encouraging sell-up to higher-priced classes leads to a 4.5% revenue gain over the base case. Although there is not a significant difference in the classes in which business passengers book, there are nearly twice as many leisure traveler bookings in E1 with Hybrid Forecasting than there are with Standard Forecasting. However, as discussed earlier, there remains an additional revenue gain that can be obtained with appropriate RM methods designed specifically for fare family structures, namely, QFF.

In all symmetric RM simulations, both Airline 1 and Airline 2 generate similar results (revenues, load factors, booking class mix, etc.), given that both airlines serve the same markets and operate the same number of legs. As such, for all experiments in the symmetric environment, only the performance of Airline 1 will be described. That is, only the results from Airline 1 will be presented, with the understanding that the results for Airline 2 would look nearly identical.

The maximum revenue obtained with each QFF method for Airline 1 (in the symmetric RM environment) is shown in Figure 5-3. The percentages above each bar graph show the revenue gains with each QFF method over the base case (represented by the horizontal dotted line). While QFF2 only leads to a 5.2% gain, both QFF1 and QFF3 both lead to revenue increases of above 7%.

Although QFF3 generates only a 68% load factor (which would be unacceptably

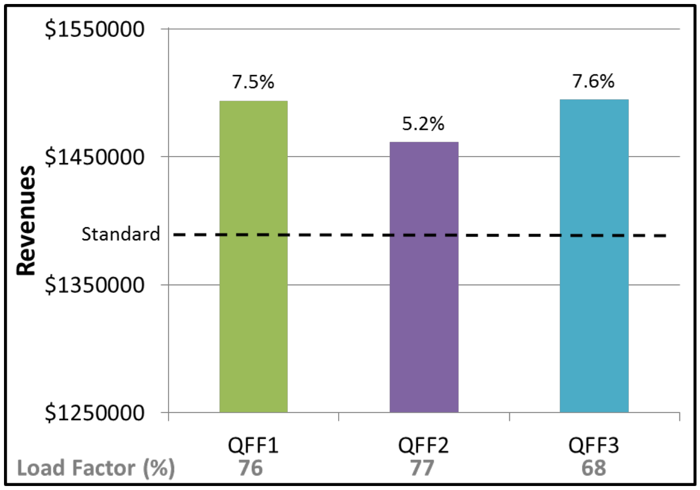

Figure 5-3: Revenues with QFF

low by many industry standards), it is worth re-iterating that the environment in which this load factor is achieved is not practical in the real world. However, testing each QFF method under such circumstances allows for a more appropriate analysis of the validity of each algorithm (without any competitive feedback). Thus, although a low load factor is obtained with QFF3, this forecasting method is performing as expected, given the symmetry of the environment. An analysis of the inputs that led to the best performance for each method, as well as the sensitivity of these parameters, is conducted in the next section.

5.1.1.1 Sensitivity Analysis of Sell-up and Buy-across Inputs

The experiments with QFF1 required inputs to model both sell-up (FRAT5) and buy-across (PBUP). Multiple combinations of FRAT5 and PBUP inputs were tested in this research, with the highest revenues being achieved with FRAT5a and a PBUP of 3.5. That is, both aggressive sell-up estimates, as well as fairly large estimates for the disutility costs attributed to the family 2 restrictions, led to the best performance for QFF1. To illustrate the sensitivity surrounding these estimates, three different FRAT5's (FRAT5a, FRAT5c, FRAT5e) were tested in combination with three separate PBUP values (4.0, 3.5, 3.0).

The revenues for Airline 1 with QFF1 with each combination of inputs are given in

Figure 5-4. More revenue is achieved for Airline 1 with more aggressive FRAT5 curves. That is, when both airlines implement aggressive sell-up inputs, many passengers are unable to book in their desired class (E2 or E1) with either airline (given that both airlines are closing classes at a similar rate). As a result, while some passengers are no-go, many others sell up to higher classes in each family, thus increasing the yield. As expected, the load factors are lower with more-aggressive FRAT5 curves. In these experiments, the network load factors range from 76% (FRAT5a) to 86% (FRAT5e).

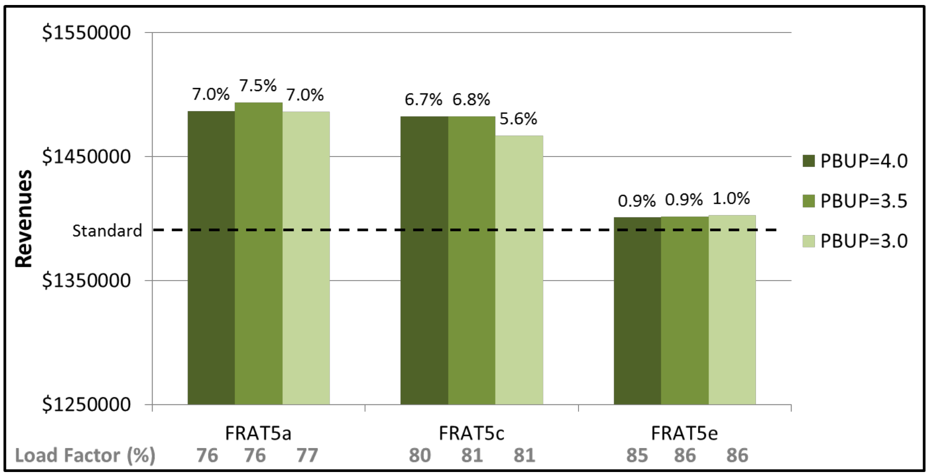

Figure 5-4: Revenues with QFF1 with different FRAT5 Curves and PBUP Values

Figure 5-4 also demonstrates the noticeable difference in impact the FRAT5 curves and PBUP values have on total revenue. While overall performance is substantially influenced with the different FRAT5 curves, the different PBUP value minimally affected the final results. Recall from Chapter 3 that this finding was one of the main motivating factors that led to the development of QFF2.

To better understand how the different FRAT5 curves impact the overall performance, the booking class mix for the three different sell-up estimates is shown in Figure 5-5 (fixed PBUP of 3.5). The classes are labeled from top-to-bottom. As expected, the aggressive FRAT5 curves inform the RM system that more passengers will sell up, which results in earlier closures for the lowest classes in each family, and thus fewer overall bookings in these classes. However, this also leads to more bookings in the top classes in both families, increasing the yield. Alternatively, with less aggressive FRAT5 curves, more bookings are made in the lower classes, resulting in an overall lower yield. Given that FRAT5a led to the highest revenues, the increase in yield outweighs the lower load factor in this case.

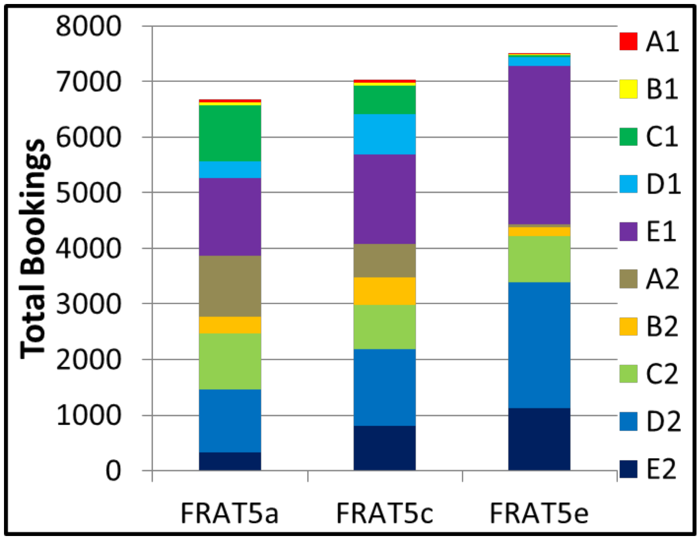

Figure 5-5: Booking Class Mix with QFF1 (PBUP 3.5) with different FRAT5 Curves

With QFF2, a separate forecast is generated for each family, thus requiring separate FRAT5 values to estimate sell-up within each family. Additionally, the buy-across input (PBUP) is not required, given that buy-up is assumed to be implicitly accounted for by the separate Q-Forecasts for each family. Consequently, the only required parameters with QFF2 are the two FRAT5 values. One of the fundamental assumptions with different forecasts for each family is that the sell-up is assumed to be constant by time frame over the entire booking period (see Section 3.3.2).

QFF2 was tested with different sell-up inputs for each fare family. Recall that the FRAT5 estimate for each family is multiplied by the lowest class in the corresponding family ($300 for E1, $100 for E2) to generate an estimate of the median WTP for passengers from each family. The highest revenues in the symmetric RM experiments with QFF2 were achieved with a family 1 FRAT5 value of 1.5 and a family 2 FRAT5 value of 2.0 (denoted as FRAT5 1.5|2.0). The sensitivity around these FRAT5 values was examined by considering different combinations of FRAT5 inputs in each family. Specifically, family 1 FRAT5 values of 2.0, 1.5, and 1.0 (median WTP estimates of $600, $450, and $300), and family 2 FRAT5 values of 2.5, 2.0, and 1.5 ($250, $200, $150) were all tested. Note that, in the case of a FRAT5 input of 1.0, no sell-up is expected, and thus the lowest family 1 class would remain open for essentially all time frames.

Figure 5-6 shows the revenues for Airline 1 with QFF2 for the nine different combinations of FRAT5 values tested. While QFF2 can lead to small gains in revenue over Standard Forecasting, it is apparent that the overall performance of the algorithm is very sensitive to the choice of the sell-up input in each family. When aggressive family 1 sell-up inputs are used in conjunction with less-aggressive family 2 sell-up inputs, revenues plummet to values *below* Standard Forecasting, as was the case with FRAT5 2.0|1.5.

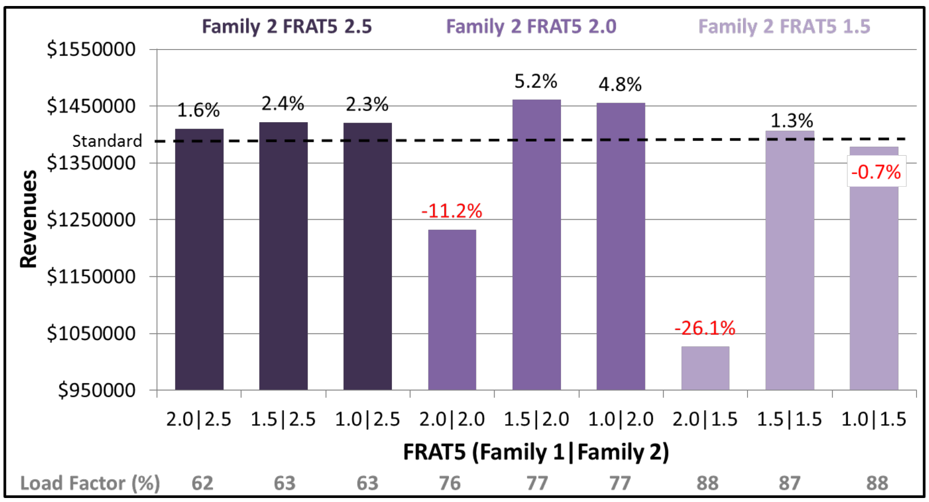

Figure 5-6: Revenues with QFF2 with different FRAT5 Values by Fare Family

As the figure illustrates, the load factor is significantly impacted by the choice of the family 2 FRAT5 input with QFF2. While a family 2 FRAT5 input of 2.5 generates load factors between 62-63%, lowering the family 2 FRAT5 to 1.5 results in load factors approaching 90%. However, the highest revenue (5.2% above the base case) was obtained with a load factor of 77%, where the family 2 FRAT5 was 2.0.

The corresponding booking class mix with QFF2 is shown in Figure 5-7 for each of the different combinations of FRAT5 values. While the family 2 FRAT5 value essentially dictates the *quantity* of overall bookings, the mix of bookings varies dramatically with different family 1 sell-up inputs. For example, with a family 1 FRAT5 of 2.0, no bookings are accepted in the lowest family 1 class. This implies that a sell-up estimate of $600 results in the RM system never opening E1 over the entire duration of the booking process. While some passengers from the family 1 forecast sell up, many are no-go. With less aggressive family 1 FRAT5 values (1.5 and 1.0), a substantial amount of bookings are made in class E1, with fewer occurring in the top family 1 classes.

The figure also demonstrates why an increase in load factor does not necessarily lead to positive results. While the load factor noticeably increases with a family 2 FRAT5 of 1.5, this is a result of an abundance of bookings being made in the lowest two classes overall. That is, very little sell-up is achieved, and, as a result of the low yield, low revenues are observed when sell-up is underestimated in family 2.

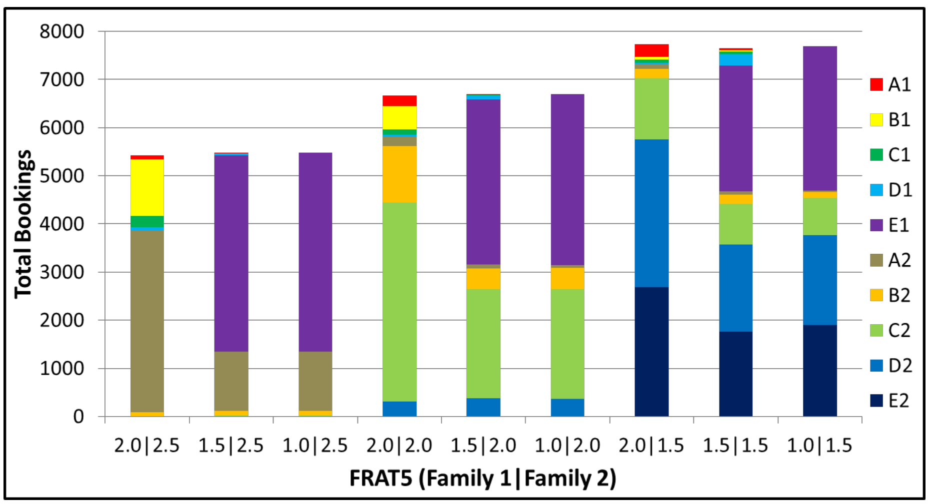

Figure 5-7: Booking Class Mix with QFF2 with different FRAT5 Inputs by Fare Family

As Figure 5-7 showed, no bookings were made in E1 with a family 1 FRAT5 of 2.0. A further analysis on the closure rates of the lowest family 1 class explains why this was the case. Figure 5-8 shows the percentage of all Airline 1 paths over the network in which class E1 was closed down for each time frames. The closure rates are shown for the three different family 1 FRAT5 values tested (with a fixed family 2 FRAT5 value of 2.0). With a family 1 FRAT5 of 2.0, QFF2 prevents E1 from becoming available for the duration of the booking process, thus preventing any bookings from being made. With the less aggressive family 1 sell-up inputs, E1 remains open for a much larger percentage of time. This indicates that, with constant FRAT5 values by time frame, there is some value for the WTP estimate for family 1 passengers that, when reached, will result in E1 never opening up throughout the entire booking process.

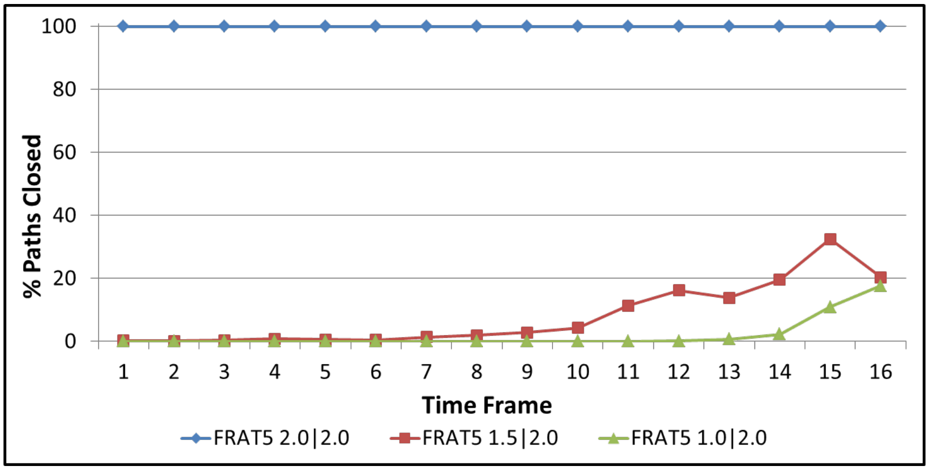

Figure 5-8: Class E1 Closure Rates with QFF2 with different Family 1 FRAT5 Values (Family 2 FRAT5 2.0)

Proceeding to the analysis of the QFF3 performance, recall that all experiments with this method required both two FRAT5 values (one for each family to estimate passengers' WTP) as well as a single DUMLT input and disutility k-factor (DUKF) estimate to model buy-across. Through multiple experiments, it was found that the same sell-up inputs that led to the highest revenue with QFF2 also led to the best performance with QFF3 (FRAT5 1.5|2.0). DUMLT inputs of 2.5, 2.0, and 1.5 were then tested on this set of FRAT5 values. It was determined that the DUKF had an insignificant impact on overall performance. As such, a fixed value of 0.5 is used in all experiments.

The revenues with QFF3 with different DUMLT inputs (assuming FRAT5 1.5|2.0) are shown in Figure 5-9(a). In general, it was determined that higher DUMLT inputs resulted in slightly higher load factors (2-3 percentage points), but overall did not have a significant impact on total revenues. A DUMLT of 2.0 (with FRAT5 1.5|2.0) resulted in a 7.6% increase in revenue over the base case. Overall, the load factors were relatively low with QFF3, with the highest revenues attaining only a 68% occupancy rate.

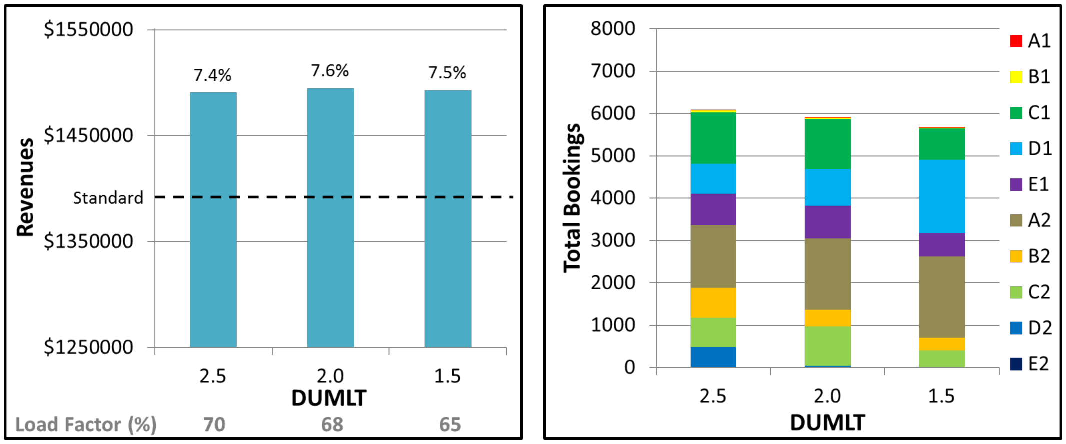

(a) Revenues (b) Bookings

Figure 5-9: Results with QFF3 (FRAT5 1.5|2.0) with different DUMLT Inputs

Figure 5-9(b) shows the booking mix for Airline 1 for the three DUMLT inputs mentioned above. More bookings occur in the top family 1 classes, and in family 2 overall, with larger DUMLT inputs. As described in Chapter 3, higher DUMLT estimates imply that the restrictions attributed to the family 2 classes will have a larger impact on passengers (from both family forecasts). When this is the case, the RM system then assumes that relatively few bookings will occur in the bottom classes (due to the unattractive restrictions), and thus keeps the lower family 2 classes open longer. This allows for more bookings in these classes, which is what was observed in the figure.

Alternatively, a low DUMLT signals to the RM system that the family 2 restrictions do not significantly impact passengers, implying that more passengers will purchase tickets in the lowest classes. As a result, the RM system closes these lower classes down early, thus resulting in fewer bookings. Regardless, the DUMLT input is found to have a much smaller influence on the performance of QFF3 than either of the FRAT5 inputs.

Throughout this section it has been shown that the success of QFF is largely dependent upon the appropriate choice of sell-up and buy-across parameters. Each method was also tested with DAVN in the symmetric environment; the results are shown in the following section.

5.1.2 OD Controls in Symmetric RM Experiments

In the previous section all of the results with the different QFF methods assumed a leg-based EMSRb seat allocation model. In practice, many of the world's largest airlines have moved towards OD control mechanisms; in this research, DAVN is used as the optimizer in all OD RM experiments. The results from the simulations in which both airlines implemented QFF with DAVN are shown in this section. These results are measured against a "new" base case, where both airlines use DAVN with path/class forecasting and an AP restriction on the different classes. In this new base case, revenues for Airline 1 were 1.1% higher than in the base case with EMSRb, a typical gain for OD RM over leg-based RM.

The revenues with OD controls in a symmetric RM environment with Standard Forecasting, as well as with each of the QFF methods, are shown in Figure 5-10. Similar trends in the revenues as from the leg-based simulations are observed. Specifically, QFF1 and QFF3 generate significant gains in revenue over the base case, while the gain with QFF2 is much lower. The load factors have increased between 1-3 percentage points for each of the methods, contributing to the overall higher revenues achieved with DAVN. Although the percentage increase in revenue due to QFF is similar with DAVN as it is with EMSRb, on the absolute scale DAVN is generating higher revenues than EMSRb, given that the revenues are higher in the OD RM base case.

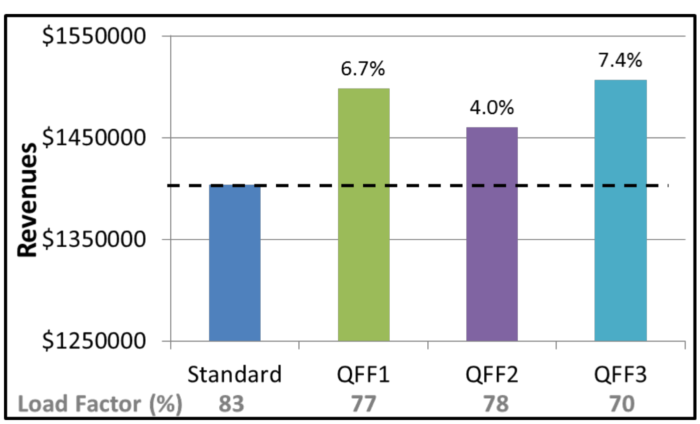

Figure 5-10: Revenues with Standard Forecasting and QFF

The sensitivity analysis of the QFF inputs in OD RM showed similar results as to the leg-based RM experiments and is not shown here. Although multiple parameters were tested with each method, the same inputs that led to the highest revenues in the symmetric leg-based RM experiments also led to the to the highest revenues with DAVN. These parameters are shown in Table 5.2.

Environment	AL1 Optimizer	AL1 Forecaster	Sell-up	Buy-across
Symmetric RM	EMSRb	QFF1	FRAT5a	PBUP 3.5
		QFF2	FRAT5 1.5|2.0	N/A
		QFF3	FRAT5 1.5|2.0	DUMLT 2.0
	DAVN	QFF1	FRAT5a	PBUP 3.5
		QFF2	FRAT5 1.5|2.0	N/A
		QFF3	FRAT5 1.5|2.0	DUMLT 1.5

Table 5.2: QFF Inputs used by Airline 1

5.1.3 Leg-based Controls in Competitive RM Experiments

Given that most airline networks in today's industry are competitive, successful performances by the different RM methods in competitive scenarios is essential. This section shows the leg-based RM results from the experiments that were performed in the competitive RM environment. Both airlines use EMSRb, with Airline 2 remaining with Standard Forecasting (with AP) throughout all runs while Airline 1's forecasting and optimization methods are modified. The same experiment that was classified as the base case in the symmetric setup is also used as the base case in the competitive environment. That is, the experiment in which both airlines use EMSRb with Standard Forecasting and an AP requirement.

In addition to testing each of the QFF methods, Hybrid Forecasting with FRAT5c (and an AP requirement) was also tested with Airline 1 in the competitive RM environment. The revenues and load factors for both airlines in this setup are shown in Figure 5-11. The forecasting method Airline 1 implements in each case is shown below the corresponding bar graphs. QFF leads to large revenue gains for Airline 1 in the competitive cases, with QFF3 generating a 13% increase in revenue over Standard Forecasting. It is interesting that, while QFF3 produced only a 68% load factor in the symmetric test case, it generates the highest load factor out of the QFF methods in the competitive scenarios (80%). As expected, HF performs much worse compared to QFF, generating only a 4.0% gain in revenue over the base case.

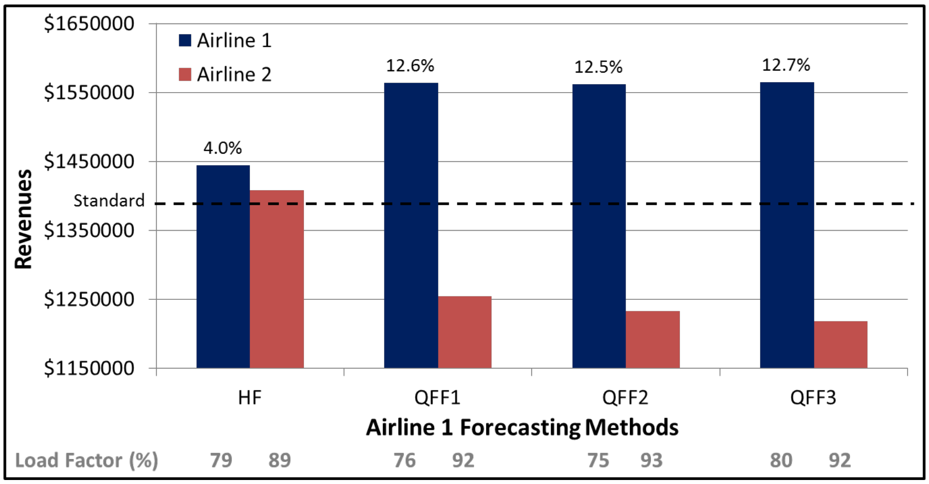

Figure 5-11: Revenues when Airline 1 uses Hybrid Forecasting and QFF

Although the focus of this analysis is primarily on the performance of Airline 1, it is interesting to observe how Airline 2 is impacted in the different competitive RM cases. For example, when Airline 1 implements HF, Airline 2 still generates a reasonable revenue at an 89% load factor. However, when Airline 1 uses QFF, Airline 2's revenue drops tremendously. Given that Airline 2's load factor actually increases in these cases (up to 93%), the lower revenues can be attributed to its substantial decrease in yield in the competitive RM experiments.

To further understand the impacts of QFF, the bookings by class are shown for Airline 1 and Airline 2 in Figure 5-12(a) and Figure 5-12(b), respectively. A quick analysis shows the drawbacks of applying HF to a fare family structure. With this method, little sell-up occurs, and a high percentage of the overall bookings are made in family 2. Alternatively, with QFF, many of the bookings are made in family 1, in large part due to the substantial amount of passengers purchasing in E1. While Airline 1 receives many passengers in the lowest family 1 class, Airline 2 accumulates more passengers in the lowest family 2 classes.

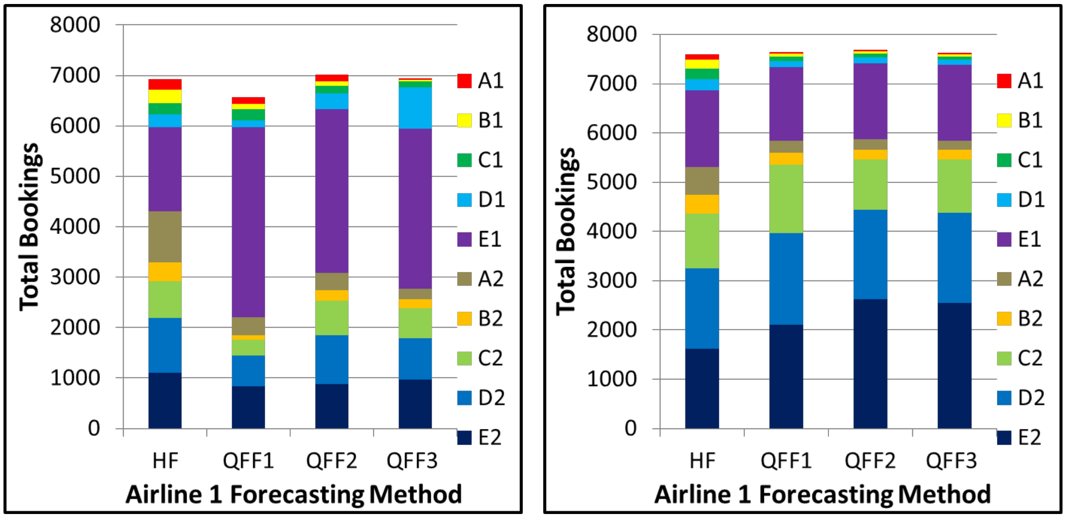

(a) Airline 1 (b) Airline 2

Figure 5-12: Booking Class Mix when Airline 1 uses HF and QFF

5.1.3.1 Sensitivity Analysis of Sell-up Inputs

Earlier it was shown that the success of QFF (in particular, QFF2 and QFF3) in the symmetric RM environment was largely dependent upon the sell-up estimates. In this section the impacts of the FRAT5 values for the QFF methods is briefly examined to determine if any major differences exist between the symmetric and competitive RM environments in regards to the sensitivity of these estimates. The analysis for QFF2 is omitted, given that the sensitivity results from the symmetric test cases with QFF2 and QFF3 were nearly identical.

Recall that simulations with QFF1 in the symmetric RM environment confirmed that the FRAT5 curve had a much larger impact on overall performance of QFF1 than the PBUP value. This conclusion remains valid in the competitive RM scenarios as well. As such, the results with QFF1 in the competitive RM environment are shown with different FRAT5 curves and a fixed PBUP input of 2.0.

The revenues for both airlines in this competitive setup where Airline 1 uses QFF1 are given in Figure 5-13. Unlike in the symmetric RM environment where both airlines were able to aggressively close down classes to encourage passengers to sell up, in the competitive RM scenarios the aggressive FRAT5 curves for Airline 1 results in passengers spilling over to the competitor, where the lowest classes in each family remain open longer. With FRAT5a (used by Airline 1), both airlines achieve similar revenues with very different load factors. While Airline 1 finishes with a 67% network load factor, Airline 2 is above 93%, a consequence of Airline 1's aggressive closure rates.

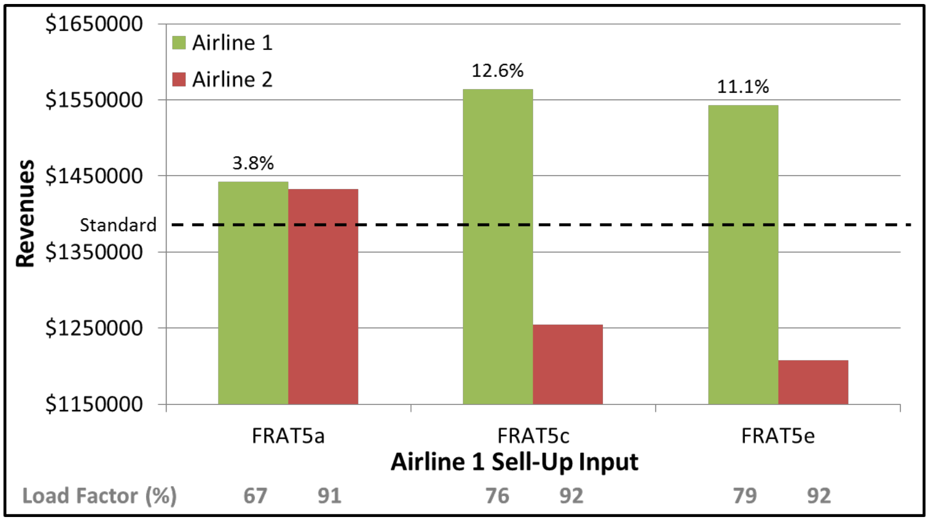

Figure 5-13: Revenues when Airline 1 uses QFF1 (PBUP 2.0) with different FRAT5 Curves

With less aggressive sell-up inputs (FRAT5c and FRAT5e), Airline 1 retains more passengers, and consequently the load factors and revenues are much higher. Although Airline 2 does not modify its RM system in any of these experiments, its revenues vary drastically depending on Airline 1's FRAT5 curves. When Airline 1 uses less aggressive sell-up estimates, revenues for Airline 2 substantially drop as a result of lower yields.

With QFF3, it was shown that in the symmetric setup the DUMLT input had only a minor impact on the overall performance of this method. In the leg-based competitive RM experiments it was determined that a DUMLT of 2.0 led to the highest revenues. However, unlike in the symmetric cases, FRAT5 1.5|2.0 performed poorly, generating revenues *below* Standard Forecasting and a load factor of only 58%. Less aggressive WTP estimates were tested by lowering the FRAT5 value in each family.

The revenues for both airlines in the competitive RM environment where Airline 1 uses QFF3 with different combinations of FRAT5 values are shown in Figure 5-14. The left three bar graphs show the revenues when Airline 1 uses a family 2 FRAT5 of 2.0, while the right three graphs depict a family 2 FRAT5 of 1.5. In the competitive RM experiments, family 1 FRAT5's were tested in 0.25 increments, given the noticeable impact that the sell-up estimates had on the overall performance.

As seen previously, aggressive family 2 FRAT5 values substantially impact the load factor. Lowering the family 2 FRAT5 to 1.5 drives up the load factor, and consequently the revenues as well. However, even more revenue can be achieved by lowering the family 1 FRAT5 to 1.25. The assumption that no family 1 passengers will sell up (family 1 FRAT5 of 1.0) was also tested. However, despite the additional increase in load factor, this did not lead to higher revenues.

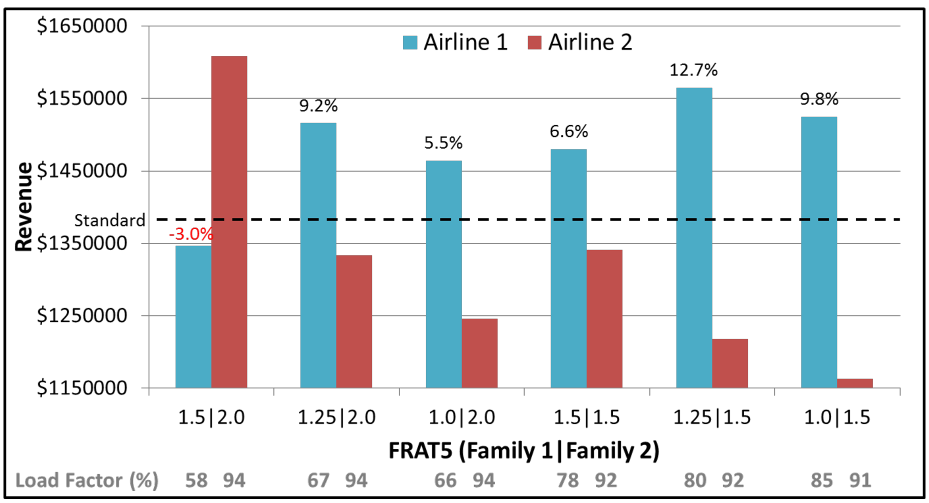

Figure 5-14: Revenues when Airline 1 uses QFF3 (DUMLT 2.0) with different FRAT5 Values by Fare Family

There are several key takeaways that are worth highlighting from Figure 5-14. First, as indicated by the wide range of load factors and revenues obtained with the different inputs, appropriate sell-up estimates are essential to the success of QFF3. Second, when a family 2 FRAT5 value of 2.0 is used, the load factor plummets (regardless of the family 1 sell-up estimate), as many passengers spill over to the competitor airline. Lastly, although less aggressive sell-up inputs lead to higher revenues, there reaches a point where the decrease in yield outweighs the higher load factor, and thus the revenues no longer increase.

While the increments of 0.25 appear to be very precise, recall that the family 1 FRAT5 value is multiplied by $300 (E1 fare). Thus in this case a 0.25 change in FRAT5 is in reality a $75 difference in the WTP estimate. To more clearly visualize the effects of the different FRAT5 values, the bookings for Airline 1 are shown in Figure 5-15. As expected, the different FRAT5 values from each family lead to very different mixes. When a family 1 FRAT5 value of 1.0 is used, all family 1 bookings are made in E1.

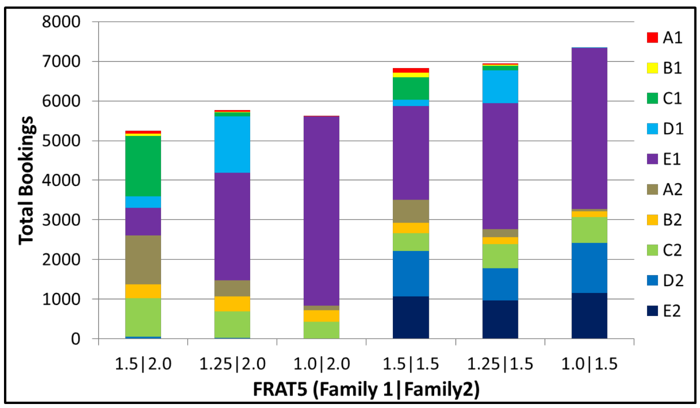

Figure 5-15: Airline 1 Bookings with QFF3 (DUMLT 2.0) with different FRAT5 Values by Fare Family

It is interesting to observe how the FRAT5 value from a particular family impacts the passenger mix in the opposite family. For example, with FRAT5 1.5|2.0, the majority of family 1 bookings are made in C1, with relatively few occurring in the lowest family 1 class. However, with FRAT5 1.5|1.5, even though the family 1 WTP estimate has not changed, the family 1 passenger mix is substantially different, with a high quantity of passenger bookings in E1. This is a consequence of the fact that the more aggressive sell-up estimates in one family lead to earlier lower-class closures, and thus indirectly influence the booking limits in the opposite family. This finding is seen frequently throughout this thesis, where the booking mix for family 1 is notable affected by the family 2 FRAT5 value, and vice versa.

5.1.4 OD Controls in Competitive RM Experiments

The QFF methods for fare family structures were developed under the assumption of an OD control mechanism. As such, a detailed analysis of the performance of QFF is provided in this section, with the focus primarily centered on QFF3.

The experiments in this section include cases in which Airline 1 uses a network seat allocation model (DAVN) with different QFF methods, while Airline 2 remains with EMSRb with Standard Forecasting and an AP requirement. As an initial base run, Airline 1 was given DAVN with Standard Forecasting (with AP).

The revenues with QFF in OD RM are shown in Figure 5-16. Airline 1 achieves higher load factors in the competitive case with DAVN than with EMSRb, and substantial increases in revenue over the base case overall. Each QFF method leads to revenue gains of approximately 14% over Standard Forecasting. The revenues for the competitor drop considerably when Airline 1 uses QFF with OD controls, again consistent with what was found when Airline 1 used a leg-based model.

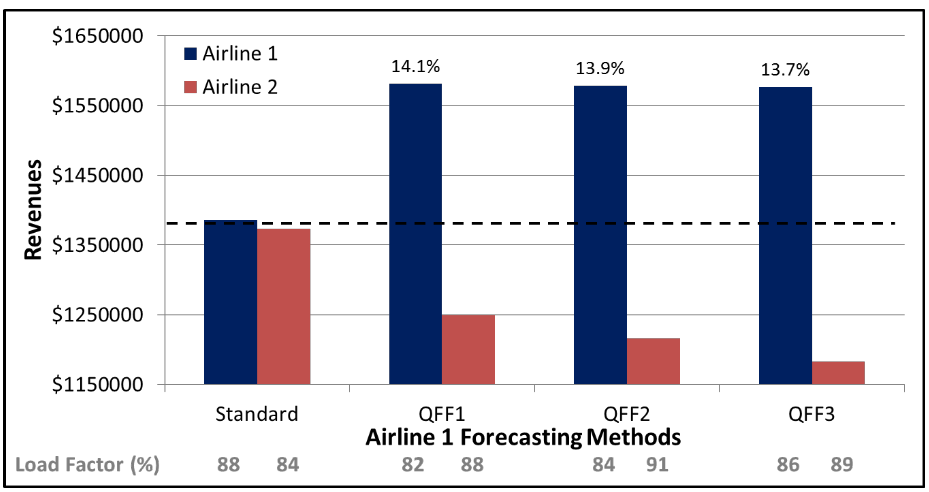

Figure 5-16: Revenues with Standard Forecasting and QFF

The same sell-up inputs that led to the highest revenues in the leg-based competitive RM simulations also led to the best performances in the OD RM experiments. Although there were slight differences in buy-across inputs, for the most part the parameters for each method were identical in both cases. These inputs are shown in Table 5.3.

Environment	AL1 Optimizer	AL1 Forecaster	Sell-up	Buy-across
Competitive RM	EMSRb	QFF1	FRAT5c	PBUP 2.0
		QFF2	FRAT5 1.5|1.75	N/A
		QFF3	FRAT5 1.25|1.5	DUMLT 2.0
	DAVN	QFF1	FRAT5c	PBUP 2.5
		QFF2	FRAT5 1.5|1.75	N/A
		QFF3	FRAT5 1.25|1.5	DUMLT 1.5

Table 5.3: QFF Inputs used by Airline 1

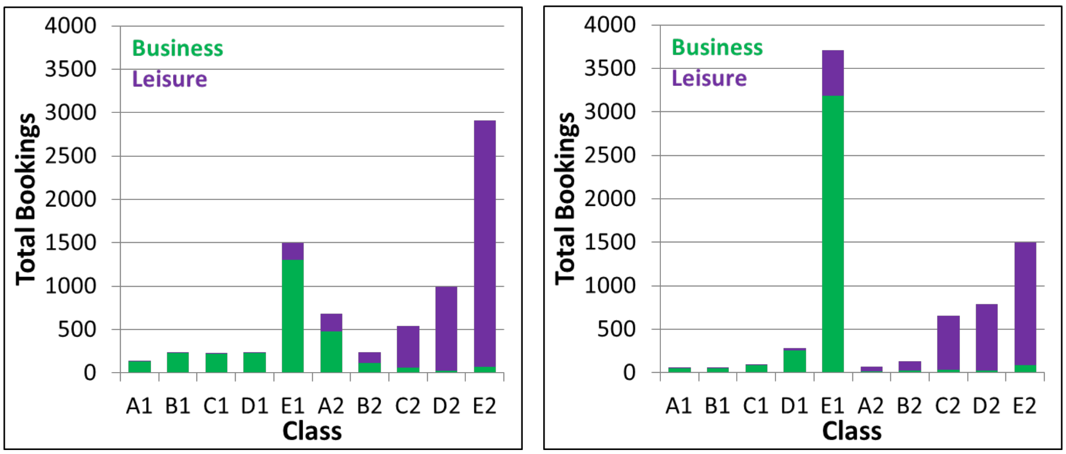

(a) Standard Forecasting

(b) QFF3

Figure 5-17: Business/Leisure Bookings for Airline 1 with Standard Forecasting and QFF3

To further explore the results with QFF in the competitive RM environment, the the business/leisure bookings for both Standard Forecasting as well as QFF3 are shown in Figure 5-17(a) and Figure 5-17(b), respectively. With the AP enforced, Airline 1 actually achieves more bookings in the top four classes with Standard Forecasting than with QFF3. However, with QFF3, E1 remains open for a much longer period of time, and thus more bookings are made in the lowest family 1 class.

While this may cause concern as to why there is more of a balanced mix with Standard Forecasting than with QFF3, recall that it is revenues, and not passengers, that is being maximized. Given that the competitor uses Standard Forecasting with AP (and thus closes E1 a minimum of 14 days before departure), the QFF3 algorithm attempts to generate the highest revenues by keeping E1 open for a longer period of time, which results in a high quantity of bookings being made in this class. A different set of RM methods for the competitor airline would likely result in a different mix for Airline 1 than what is seen in Figure 5-17(b).

It is worth noting that there are approximately 33% more business traveler bookings with QFF3 than with Standard Forecasting. Fewer travelers spill over to the competitor, and less no-go is observed. Also, as a consequence of E1 remaining open for a longer period of time, there are more leisure passengers buying up to family 1, as well as fewer fewer business travelers buying down to family 2.

To further scrutinize the differences between Standard Forecasting and QFF3, Figure 5-18 shows the cumulative bookings by family for both Standard Forecasting and QFF3 across the different time frames. As already shown, more family 2 bookings occur with Standard Forecasting. The difference between the quantity of family 2 bookings with Standard Forecasting and QFF3 becomes larger as the departure date nears.

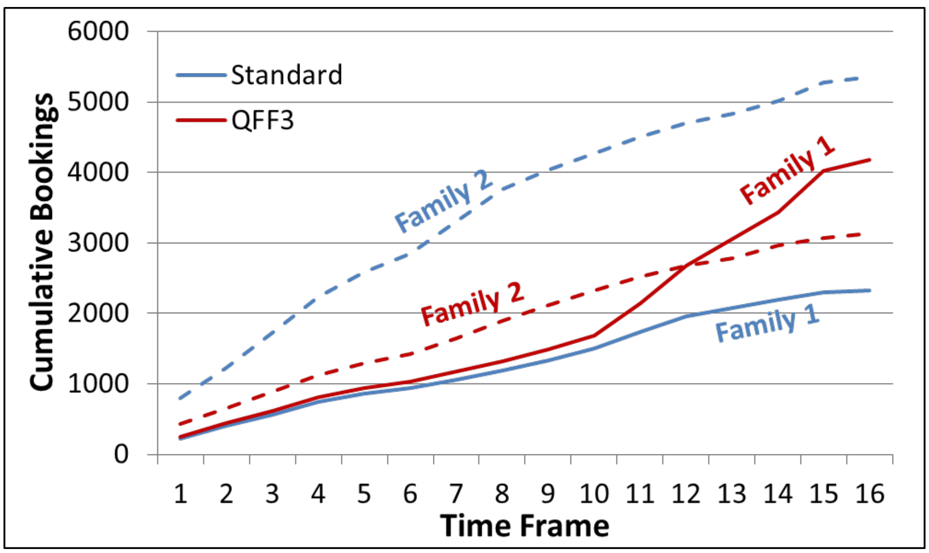

Figure 5-18: Cumulative Family 1 and Family 2 Bookings by Time Frame with Standard Forecasting and QFF3

As Figure 5-18 illustrates, the total number of family 1 bookings that each method generates is essentially the same during the beginning of the booking process. However, after about Time Frame 10, while the increase in cumulative family 1 bookings with Standard Forecasting remains nearly constant, the number of family 1 bookings increases dramatically with QFF3, owing to the number of both business and leisure travelers booking in E1. As mentioned previously, although the ratio of family 1/family 2 passengers may be different depending on the choice of RM methods for the competitor, in this particular setup, there is a higher quantity of family 1 bookings for airline 1.

5.1.4.1 Performance of QFF under Low Demand Scenarios

As mentioned in Chapter 4, PODS allows users to adjust the demand over the entire network. All of the test cases thus far have been conducted with "medium" demand (approximately 85% load factor in the base cases). Before proceeding to the results from the simulations that with overlapping fare family structures, the QFF methods are tested in a competitive environment under "low" demand settings (with non-overlapping structures) to determine how the performance of each algorithm is impacted.

In the base case with lower demand, Airline 1 generated approximately 10% lower revenues than what it obtained in the base case with medium demand. However, when the QFF methods were tested under the new settings, each of them produced similar percentage increases in revenue over Standard Forecasting as to what was achieved in the medium-demand experiments. The revenues for both airlines are given in Figure 5-19. The same QFF parameters that were used with DAVN in the medium-demand scenarios (Table 5.3) were found to generate the best results in the low-demand cases as well.

While the load factors (for Airline 1) drop approximately 2 - 4 percentage points in the lower-demand simulations, it is the yield that is especially impacted in these new runs. With fewer bookings being made, the bottom classes remain open longer, and as a result, a much lower yield is achieved. However, the relative performance of each method is nearly identical in the lower-demand experiments, with revenue gains well above 13% with QFF3. Interestingly, the load factors for Airline 2 are now between 83% to 85%, much more realistic than the 93% that was achieved in the medium-demand experiments.

To more clearly see the impacts the medium-demand and low-demand situations have on the different passengers, the business/leisure mix for Airline 1 with QFF1

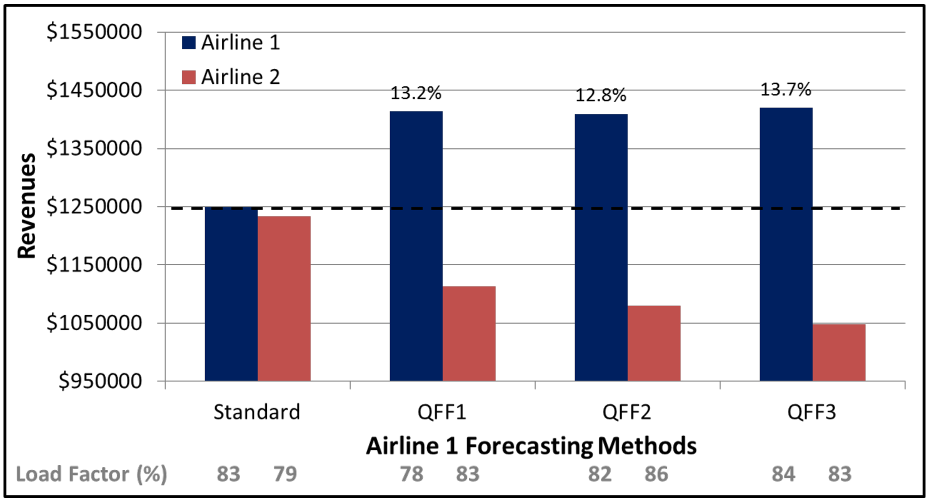

Figure 5-19: Revenues in Low Demand Experiments when Airline 1 uses Standard Forecasting and QFF

in the competitive environment is shown in Figure 5-20(a) (medium demand) and Figure 5-20(b) (low demand). Both mixes are from experiments in which FRAT5c with PBUP 2.5 were used. The yield is much smaller in the low demand setting, with fewer bookings in family 2. Overall, there are nearly half as many bookings in E1 with lower demand, while almost twice as many in E2.

Many of these experiments performed with a non-overlapping structure were also conducted with an overlapping fare family structure in order to determine how the relative performance of each method was impacted. These results are presented in the following section.

5.2 Overlapping Fare Family Structures

The fare family structures with overlapping price points included higher fares in the top family 2 classes (compared to the non-overlapping structure), as well as lower fares in the bottom family 1 classes. As the results are presented, it is worth noting that all of the experiments in this section are completely independent from the simulations in which non-overlapping structures were used, given the major differences between the prices of the fares. As such, for the remainder of this chapter, all of the results mentioned will refer to the experiments performed when both airlines used the

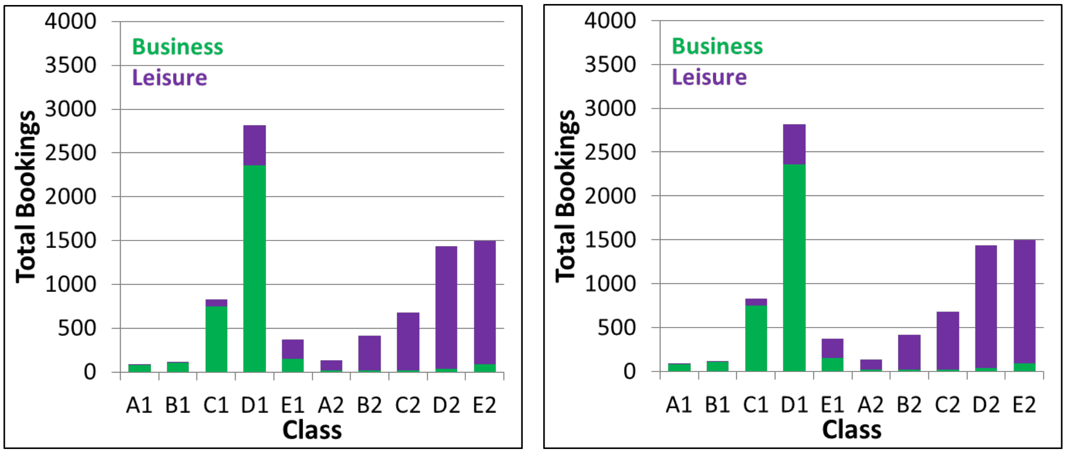

(a) Medium Demand (b) Low Demand

Figure 5-20: Airline 1 Bookings with QFF1 with different Demand Levels

overlapping structure shown in Figure 4-7. As in the previous section, this section will begin by showing leg-based and network results in the symmetric RM test cases, and then culminate with results in the competitive scenarios.

5.2.1 Symmetric RM Experiments

In this subsection the results from the symmetric RM test cases, both with EMSRb and DAVN, are presented. An analysis of the sensitivity of the FRAT5 values with DAVN is then covered in order to determine if the impact of the sell-up estimates changes from one type of fare family structure to another.

The revenues for Airline 1 with both EMSRb and DAVN are given in Figure 5-21. The leftmost bar graphs include the revenues when both airlines use Standard Forecasting. The percentage increases in revenue with EMSRb and QFF are all compared to EMSRb with Standard Forecasting, while the increases with DAVN are measured against DAVN with Standard Forecasting.

The revenue increases with QFF are more substantial in the overlapping structure than they were in the non-overlapping structure, particularly with QFF1 and QFF3. Each of these methods generates approximately 10% revenue increases over the respective base cases. As previously observed, QFF2 produces significantly lower gains in revenue as compared to the other methods. With DAVN, each QFF method

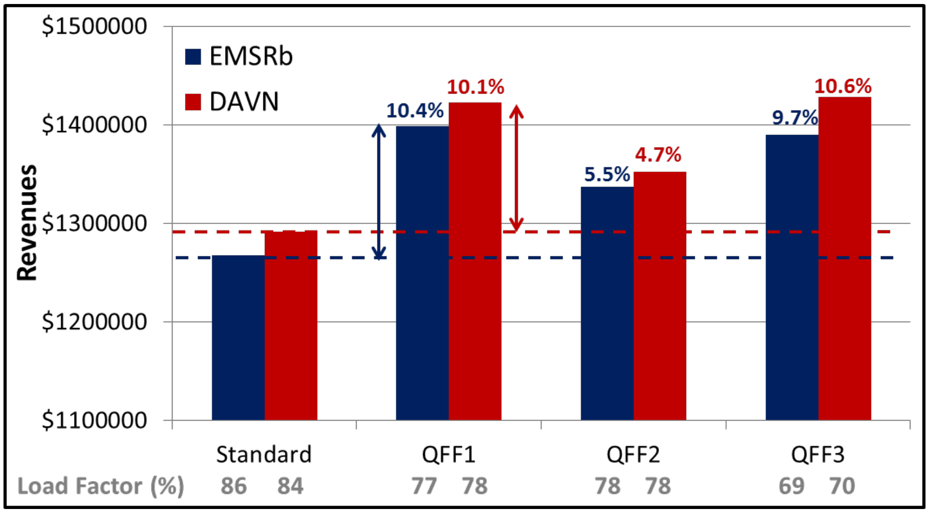

Figure 5-21: Revenues with Standard Forecasting and QFF in both Leg-based and OD RM Simulations

makes an even larger impact on overall performance, further increasing the revenues. Although the percentage increase with QFF1 and QFF2 with DAVN is less than with EMSRb, this is a result of the base case with DAVN generating 2% higher revenues than in the EMSRb case case.

The inputs that led to the highest revenues were nearly identical with both EMSRb and DAVN, and are shown in Table 5.4. While the QFF1 parameters are very similar as to what they were with the non-overlapping structure, the family 1 FRAT5 values with QFF3 are much higher than before . However, recall that the family 1 FRAT5 value is multiplied by the lowest family 1 class to estimate the median WTP among passengers from the family 1 forecast. Given that the E1 fare is now $200 (previously $300), the estimate for the median WTP is not far off from what it was in the previous fare structure. The sensitivity of the QFF3 FRAT5 values (with DAVN) is explored in the following section.

Environment	AL1 Optimizer	AL1 Forecaster	Sell-up	Buy-Across
Symmetric RM	EMSRb	QFF1	FRAT5a	PBUP 3.0
		QFF2	FRAT5 2.0|2.0	N/A
		QFF3	FRAT5 3.0|2.0	DUMLT 1.5
	DAVN	QFF1	FRAT5a	PBUP 3.5
		QFF2	FRAT5 2.0|2.0	N/A
		QFF3	FRAT5 3.0|2.0	DUMLT 1.5

Table 5.4: QFF Inputs used by Airline 1

5.2.1.1 Sensitivity of Sell-Up Inputs with QFF3 with OD Controls

Before analyzing the QFF3 inputs, the booking class mix with DAVN is shown in Figure 5-22. In the symmetric RM test cases, QFF1 achieves bookings in all 10 classes, while both QFF2 and QFF3 are more aggressive in the lowest classes in each family. These traits were also observed in the non-overlapping structure.

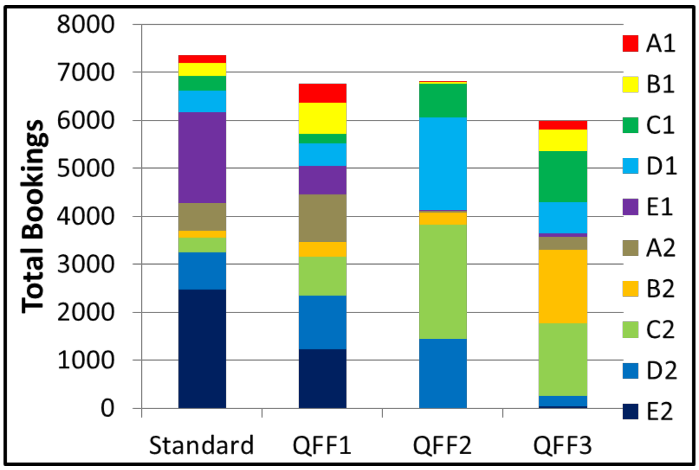

Figure 5-22: Bookings Class Mix with DAVN

As Table 5.4 indicated, FRAT5 3.0|2.0 was used with QFF3 in the symmetric DAVN experiments. When other FRAT5 values were tested, the family 2 input largely impacted the overall performance (as seen in previous results). Given the reduced E1 fare, this was even more evident with Overlapping fares.

The revenues with QFF3 (DUMLT 1.5) and different FRAT5 values are presented in Figure 5-23. With an even higher family 2 FRAT5 value (2.5), revenues actually increase. However, the load factor decreases to around 60%. Alternatively, lowering the family 2 FRAT5 to 1.5 brings the load factor up, but the decrease in yield is more substantial, thus leading minimal revenue gains above Standard Forecasting. This result confirms that, in a symmetric RM environment, revenues are largely dictated by the WTP estimate for passengers in the family 2 forecast.

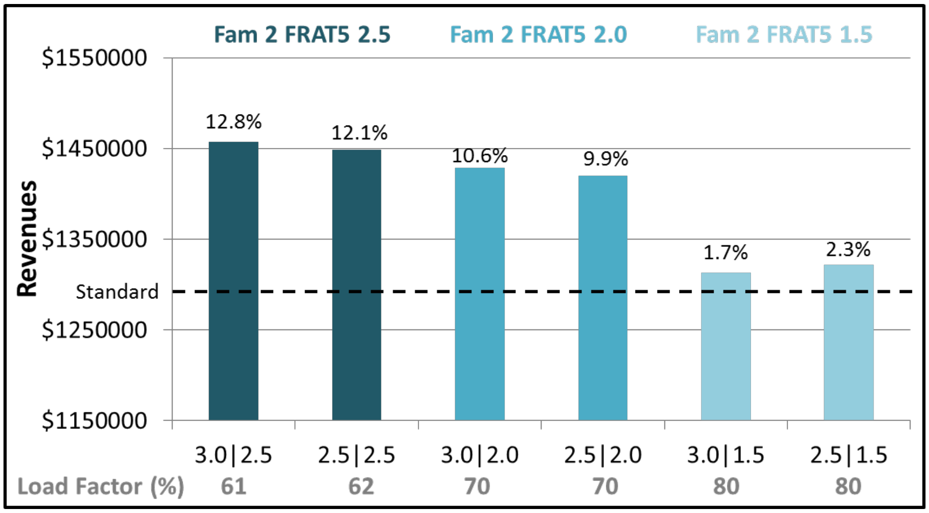

Figure 5-23: Revenues with QFF3 with different FRAT5 Values

5.2.2 Leg-based Controls in Competitive RM Experiments

There were several intriguing findings in the competitive RM test cases with overlapping fares. This section begins by examining the cases in which Airline 1 used EMSRb with QFF. The results throughout this section are measured against the revenue obtained with EMSRb and Standard Forecasting from Figure 5-21, that is, $1,267,698.

The revenues for both airlines when Airline 1 uses QFF are given in Figure 5-24. With QFF1, Airline 1 achieves a 7.6% gain in revenue over the base case at an 81% load factor. While QFF2 also generates an 81% load factor, only a 2.0% revenue increase is obtained, thus implying a lower yield. Finally, QFF3 leads to minimal gains over Standard Forecasting, a consequence of a 67% load factor. Note that while QFF3 consistently produced low load factors in the *symmetric* RM experiments, this was not the case in the competitive RM scenarios. It is also interesting to note that, although Airline 1 uses QFF3 and Airline 2 uses Standard Forecasting, both airlines achieve similar revenues. This oddity is investigated in the following section.

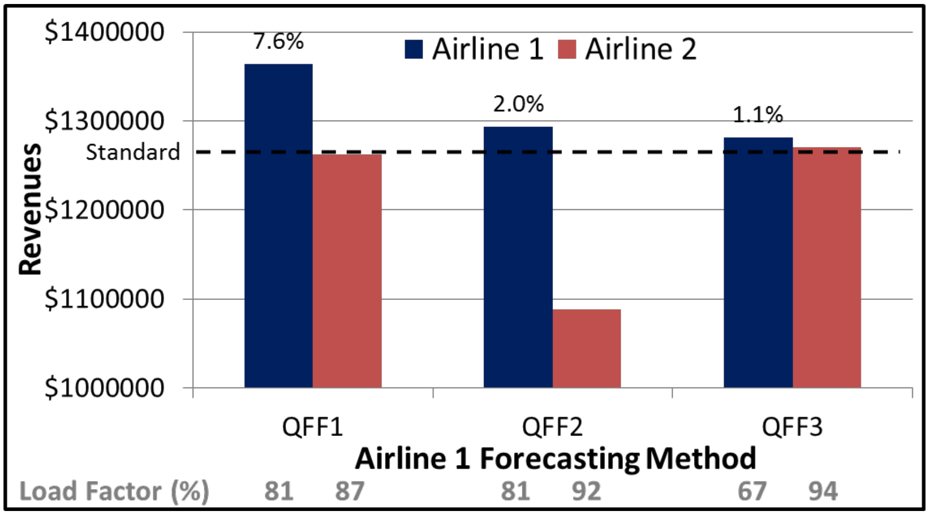

Figure 5-24: Revenues with QFF

5.2.2.1 Analysis of QFF3 Performance in Competitive RM Experiments

To gain a better insight as to the drawbacks of QFF3, the bookings are analyzed in comparison to the bookings achieved with Standard forecasting, given that both methods generated similar revenues (albeit very different load factors). The business/leisure booking mix for Standard Forecasting and QFF3 is shown in Figure 5-25(a) and Figure 5-25(b), respectively. With a reduced E1 fare, there are many more leisure passengers purchasing the lowest family 1 class with Standard Forecasting than there were in the Non-Overlapping structure. Because of this, when E1 eventually does close down, those business passengers that were displaced by the leisure passengers either are forced to sell-up, or buy-down to family 2 (very few passengers are spilled, given that the competitor is undergoing a similar process). This explains why a fair amount of business bookings are made in the top family 2 classes.

While there are slightly more business passenger bookings with QFF3, there are nearly 50% more leisure passenger bookings with Standard Forecasting, a result of more leisure passengers buying in both family 1 and family 2. With QFF3, a relatively low amount of bookings are made in the lowest class in each family. It is evident that these classes are being closed down very early, forcing passengers to make a different decision on which class they will book in, if any. While some passengers sell-up, the majority either are spilled to Airline 2 or are no-go.

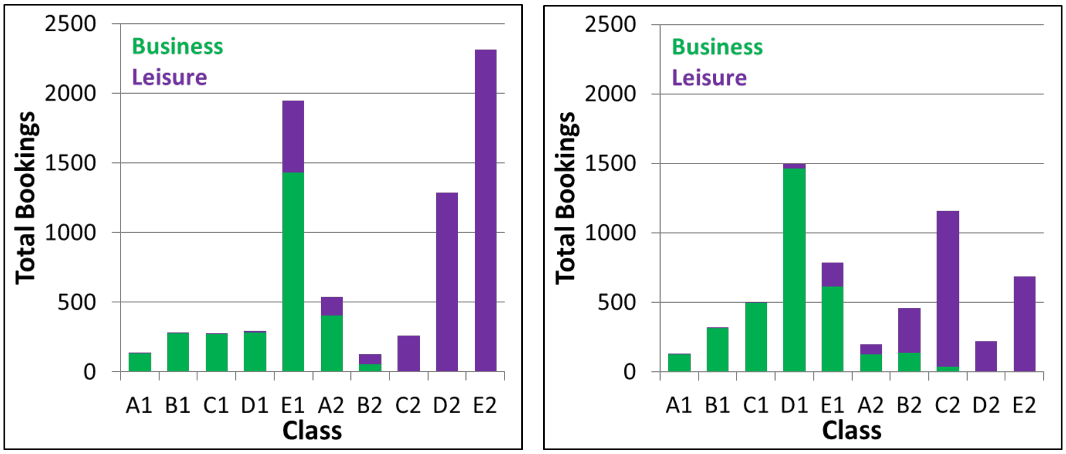

(a) Standard Forecasting (b) QFF3

Figure 5-25: Bookings for Airline 1 with Standard Forecasting and QFF3 in Competitive RM Experiments

To more clearly illustrate passenger choice, Figure 5-26 shows the details of the actual passenger decision, given a first choice of E1 or E2. Recall that, in the fare family structures used throughout this thesis, all passengers have a first choice of E1 (business travelers) or E2 (leisure travelers), given the lack of restrictions between the classes within each family. With QFF3 implemented, many passengers are unable to book in their desired class (either E1 or E2), given the early closure rates of each of these classes. Out of the approximately 4400 business passengers generated throughout this network, only a little more than 10% received their first choice (E1). While more than half of all the business passengers sell-up (captured in the "Booked in Family 1" category), a significant amount of travelers spill over to the competitor. The business travelers that bought down to family 2 are represented by the green bar graph in the "Booked in Family 2" column.

An even smaller percentage of all the leisure passengers are able to book in E1, with only about 6.3% of these travelers obtaining their first choice. As was the case with the business travelers, while some leisure passengers sell-up to higher family 2 classes ("Booked in Family 2"), an even larger percentage are either spilling over to Airline 2 or are no-go. As a result, QFF3 results in insignificant gains in revenue over Standard Forecasting alone.

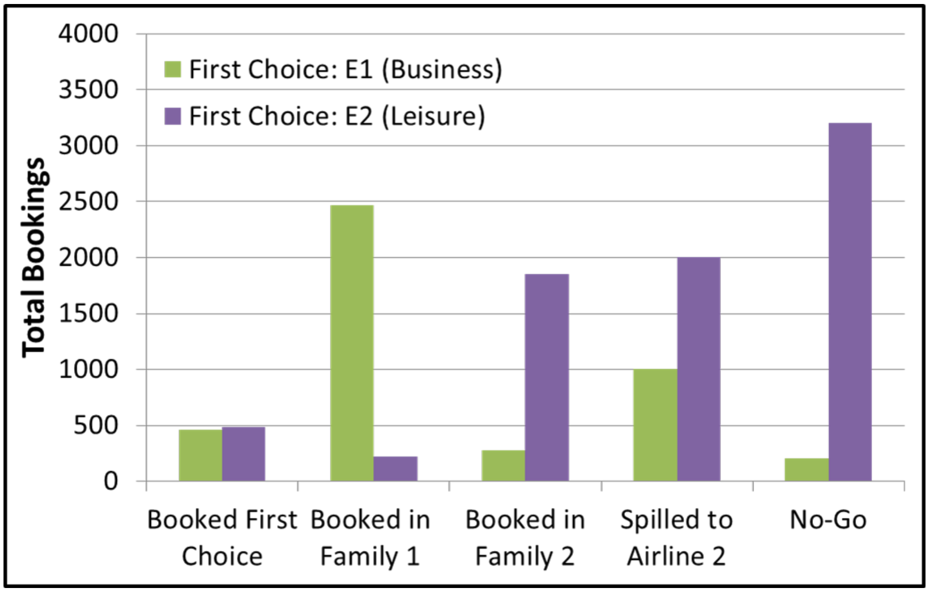

Figure 5-26: Breakdown of Passenger Decision with QFF3, given a First Choice E1 (Business) and First Choice E2 (Leisure)

Given that early class closures have been identified as the main contributor to the poor mix of bookings with QFF3, further investigations were conducted to determine the reason for these aggressive closure rates. It was discovered that the main issue with QFF3 was with the original Q-Forecasts generated for each family. As described in Chapter 3, the first step in the QFF3 algorithm is to generate Q-Forecasts by family using historical data from each family. The fundamental assumption in this step is that the majority of bookings in family 1 are from business travelers, while the majority in family 2 are from leisure passengers.

Although this assumption was reasonably valid with non-overlapping fare family structures (given the gap in fares between the two families), this was not the case with overlapping fares. More leisure passengers booked in family 1 and more business passengers booked in family 2 than with the original fare family structure. As a consequence of the way the bookings are scaled, higher Q-Forecasts were generated in each family across all the time frames. With higher Q-Forecasts, more protection was given to the top classes in each family, thus leading to aggressive closure rates in the lower classes (in each family) with QFF3.

5.2.3 OD Controls in Competitive RM Experiments

To culminate the matrix of results presented throughout this chapter, Figure 5-27 shows the revenues with DAVN in the competitive RM environment. The inputs used to generate these revenues, as well as the highest revenues in the leg-based competitive RM simulations, are given in Table 5.5. For the reasons listed above, QFF3 expectedly performs poorly compared to QFF1. While a slightly higher load factor is obtained when DAVN is implemented with QFF3, only a 2.1% increase in revenue is achieved, much lower than the 6.7% gained with QFF1.

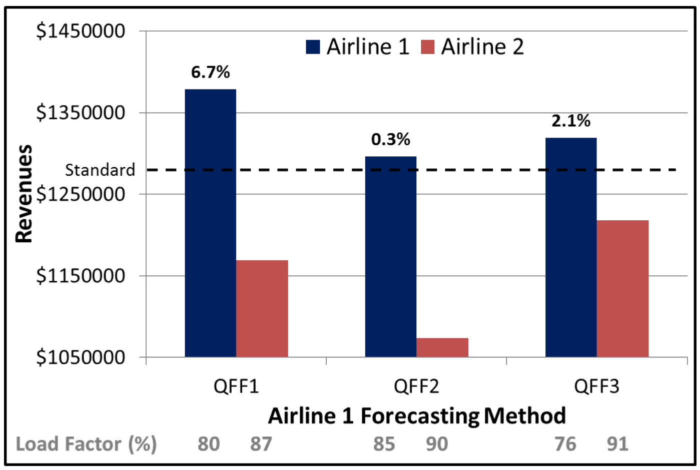

Figure 5-27: Revenues with QFF

Environment	AL1 Optimizer	AL1 Forecaster	Sell-up	Buy-across
Competitive RM	EMSRb	QFF1	FRAT5c	PBUP 2.5
		QFF2	FRAT5 1.5|1.75	N/A
		QFF3	FRAT5 2.0|1.5	DUMLT 1.5
	DAVN	QFF1	FRAT5c	PBUP 1.5
		QFF2	FRAT5 1.5|1.75	N/A
		QFF3	FRAT5 2.0|1.5	DUMLT 1.0

Table 5.5: QFF Inputs used by Airline 1

5.3 Chapter Summary

In this chapter, the results from the simulations with both non-overlapping and overlapping fare family structures were presented. With each type of fare structure, the three QFF methods were tested with both a leg-based RM optimizer (EMSRb) as well as with OD controls (DAVN). These methods were tested first in the symmetric RM environment to determine if the algorithm produced the desired results assuming no competitive feedback. These methods were then analyzed in the competitive RM environment, where Airline 2 was assigned a fixed optimization and forecasting method (EMSRb with Standard Forecasting and AP).

Beginning with simulations with non-overlapping fare family structures, the performance of QFF was measured against a base case in each subsection. The revenue gains with both EMSRb and DAVN in the different environments are summarized in Table 5.6. Regardless of the optimizer, QFF2 generated much lower revenues than QFF1 or QFF3 in the symmetric test cases. This is most likely a result of the fact that buy-across is not modeled within the algorithm. With all methods, DAVN led to higher overall revenues.

Environment	**Forecasting Method**	**Optimizer**	
		EMSR	DAVN
Symmetric RM	QFF1	7.5%	6.7%
	QFF2	5.2%	4.0%
	QFF3	7.6%	7.4%
Competitive RM	QFF1	12.6%	14.1%
	QFF2	12.5%	13.9%
	QFF3	12.7%	13.7%

Table 5.6: Summary of Revenue Gains with QFF over Standard Forecasting with Non-Overlapping Fare Family Structures

In the competitive test cases, all QFF methods generated much larger revenue gains over Standard Forecasting. This was a result of the high quantity of bookings made in the lowest family 1 class. However, as mentioned earlier, the mix of passengers for Airline 1 is largely dependent upon the choice of RM system assigned to Airline 2.

Similar experiments were then shown for the test cases with overlapping fare family structures. The results from these simulations are summarized in Table 5.7. In the symmetric environment, QFF1 and QFF3 make an even larger impact than before, with revenue gains of approximately 10% being achieved. QFF2 is once again seen to generate much lower revenues than either of the other methods in the symmetric cases.

In the competitive scenarios, it was discovered that both QFF2 and QFF3 perform poorly, posting minimal gains over Standard Forecasting. With buy-across not being modeled with QFF2, the results with this method were not completely surprising. However, the poor performance by QFF3 led to further investigations, where it was shown that the Q-Forecasts from each family were inaccurate due to high quantity of leisure travelers buying up to family 1 (and business passengers buying down to family 2). The way the bookings were scaled overestimated the forecasts in each family, thus leading to earlier lower class closures while overprotecting the highest classes.

Environment	Forecasting Method	Optimizer	
		EMSR	DAVN
Symmetric RM	QFF1	10.4%	10.1%
	QFF2	5.5%	4.7%
	QFF3	9.7%	10.6%
Competitive RM	QFF1	7.6%	6.7%
	QFF2	2.0%	0.3%
	QFF3	1.1%	2.1%

Table 5.7: Summary of Revenue Gains with QFF over Standard Forecasting with Overlapping Fare Family Structures

Chapter 6

Conclusions

With the emergence of low cost carriers in the early 2000s, many major airlines were forced to simplify their fare structures and develop alternative means to remain competitive. The strategy of offering branded fares, or "fare families", has become more popular among airlines as a way to differentiate their products and services from competitors. Legacy carriers offered branded fares to create a better value proposition to their customers by bundling ancillary services together and attributing distinct benefits to each fare family. The airlines expected that, by better promoting brand awareness and by enabling passengers to purchase a product that better met the their specific needs, revenues could ultimately be increased.

There were numerous revenue management (RM) challenges associated with implementing this type of fare structure. Many of the RM forecasting and optimization models that came about in the 1990s were developed under the assumption of a single set of fare products. That is, the third-generation RM systems implemented by many airlines were not designed for fare family structures. Thus in order to fully take advantage of branded fares, forecasting and optimization methods that appropriately model the customer choice process in a fare family structure were needed.

6.1 Summary of Thesis Objectives

The objective of this thesis was to provide a comprehensive overview of the forecasting and optimization methods developed in the MIT PODS Research Consortium specifically for fare family structures. Three variants of fare family RM methods were developed; each was constructed based on the assumed passenger decision process in a fare family structure. These methods were collectively termed Q-Forecasting for Fare Families (QFF) due to the extensive use of Q-Forecasting concepts (described in Chapter 2) in each model.

In this thesis, all fare family structures tested consisted of two families, with five classes in each family. All classes within a particular family were undifferentiated from one another, with the family 1 classes being completely unrestricted while the family 2 classes were fully restricted.

The first QFF formulation (QFF1) generates a single Q-Forecast of demand for both family 1 and family 2 combined. Because all passengers in the forecast were assumed to follow the same sell-up function, the FRAT5 curve (an estimate for WTP in PODS) was modeled in such a way as to capture the different booking tendencies between business and leisure passengers. That is, given less price sensitive business travelers tend to book closer to departure, the probability of a given passenger selling up to a higher fare was expected to gradually increase as the departure date approached. This resulted in an "S-shaped" FRAT5 curve.

Buy-across, described in Chapter 3 as the decision a passenger makes between the lowest open family 1 class and the lowest open family 2 class, was also accounted for in QFF1 with a PBUP value. To determine how "attractive" a (more expensive) family 1 product was compared to a fully-restricted (though less expensive) family 2 product, the disutility costs the passenger in question attributed to the family 2 restrictions were required as inputs. The algorithm assumes the passenger would ultimately select the most "favorable" product after considering both the fare as well as the disutility costs.

While it accounted for both passenger sell-up and buy-across, there were limitations with the QFF1 formulation. It was discovered in the simulation of QFF1 that too much emphasis was put on the FRAT5 input. That is, the sell-up estimate had a much larger impact on the overall performance of QFF1 than did the buy-across input. Additionally, airlines from the PODS Consortium preferred a more practical estimate for sell-up; the feedback from the consortium members was to estimate sell-up separately in each family.

The QFF algorithm was refined to estimate the probability of sell-up by family. In the second formulation (QFF2), differential forecasts are generated in each family (using historical data). Because the separate Q-Forecasts are based on past bookings in each family, it is assumed that the family 1 forecast mainly consists of business passengers while the family 2 forecast represents mainly leisure travelers. As such, it is assumed that the WTP of passengers from each forecast would not significantly change as the departure date neared. That is, the passengers from the family 1 forecast were assumed to have a similar WTP throughout the entire pre-departure process, as were the passengers from the family 2 forecast. Thus, QFF2 uses two constant FRAT5 values (one for each family) as sell-up inputs in place of the S-shaped curve used in QFF1.

Additionally, under the QFF2 formulation, buy-across was assumed to be implicitly taken into account, given that the forecast for each family is based on historical data available from previous bookings in each family. Consequently, there is no buy-across input in QFF2.

While the second QFF formulation provided a more practical RM method from an airline's standpoint than QFF1, the results from PODS simulations showed that some of the key assumptions in the QFF2 model were not valid. Specifically, although buy-across was assumed to be accounted for by the separate family forecasts, simulations showed that the PBUP was important after all. A significant percentage of passengers purchased a product in the opposite family from which they were forecast. This interaction between the two families was not captured in the QFF2 formulation. These results provided the necessary motivation to further modify the QFF algorithm. The third method (QFF3) was developed with the intention of incorporating the better aspects from the first two methods into a new formulation.

As was the process with QFF2, QFF3 generated separate forecasts by family, and single FRAT5 values were required to estimate the probability of sell-up within each family. However, the second and third formulations differed from one another in the fact that QFF3 accounted for buy-across. Given that two separate forecasts were generated with QFF3, two different elements of buy-across were modeled. Specifically, "buy-up" and "buy-down". Buy-up occurs when passengers from the family 2 forecast purchase a family 1 product, while buy down refers to a passenger from the family 1 forecast booking in family 2. To include the probability of both buy-up and buy-down in the model, the parameter DUMLT was added to the formulation to estimate the disutility costs each passenger attributes to the family 2 restrictions.

Overall, of the three formulations, QFF3 best reflects the the actual passenger choice process with fare families. Although implementing QFF3 requires an additional input (DUMLT), the input for estimating the probability of sell-up is more practical than with QFF1 (which required a detailed FRAT5 curve).

In Chapter 4, the PODS simulation tool used to test the three distinct QFF formulations was discussed. All simulations took place in a dual airline competitive environment. Two types of fare family structures were tested: "non-overlapping" fare structures and "overlapping" fare structures. In the non-overlapping structures, the lowest family 1 class was priced above the highest family 2 class, whereas in the overlapping structures, the price points between the lowest family 1 and highest family 2 classes were intertwined.

6.2 Summary of Results

The QFF formulations were tested in both symmetric RM and competitive RM scenarios. In the symmetric RM test cases, both airlines in the network were assumed to use identical seat allocation models, forecasting methods, and advanced purchase requirements (or lack thereof). In the competitive RM simulations, Airline 1's RM methods were modified while the forecasting and optimization methods used by Airline 2 (chosen as the competitor in the simulations) were fixed. Specifically, Airline 2 implemented a leg-based seat optimizer (EMSRb), Standard Forecasting, and an advanced purchase requirement in all competitive RM simulations.

In general, the QFF methods showed promising results in most of the experiments in which they were tested, particularly QFF1 and QFF3. To analyze and compare the results of QFF against a commonly used forecasting method, Hybrid Forecasting with traditional optimization was also tested in this thesis. With non-overlapping fare family structures, Hybrid Forecasting was shown to produce about a 4.0% increase in revenue over Standard Forecasting, both in symmetric RM and competitive RM simulations. In contrast, QFF1 and QFF3 generated revenue gains above 7.5% and 12.5% in the symmetric RM and competitive RM scenarios, respectively. The strong performance of the QFF methods was a direct result of higher yields (in many cases the load factors with QFF were lower than with Standard Forecasting). That is, by appropriately modeling the passenger decision process in a fare family structure, QFF was able to encourage passengers from each family to sell up to higher-priced classes. With network RM controls (DAVN), the revenue gains with QFF were further magnified.

With overlapping fare family structures, similar conclusions were reached about the performance of QFF in symmetric RM scenarios. Specifically, QFF1 and QFF3 generated much higher revenues than QFF2. However, in the competitive RM experiments, QFF3 resulted in lower revenues and load factors than expected. The modeling issue was determined to be with the original Q-Forecasts, which used historical data from each family. Because the lowest family 1 classes were priced below the highest family 2 classes, more leisure passengers booked in family 1, and thus the family forecasts were not appropriately segmenting the different passenger types. As a result, QFF3 generated too high of protection levels for the top classes in each family, which ultimately led to early closure rates in the lowest classes.

For a complete summary of the revenue benefits with each QFF method with both non-overlapping and overlapping fare family structures, refer to Table 5.6 and Table 5.7, respectively.

6.3 Potential Directions for Future Research

Based on the findings in this thesis, two potential areas for research are suggested. The first direction is on the validity of the sell-up estimate used with each QFF method. One of the general trends observed throughout this thesis was that minor adjustments in sell-up inputs to the algorithms had major impacts on the revenues and booking class mix in many simulations. In some cases, QFF performed worse than Standard Forecasting. This occurred when the estimate for sell-up was too aggressive (lowest classes closed too early) or when the sell-up potential was under estimated (lowest classes left open too long). In either case, a poor sell-up estimate had dramatic effects on the overall performance of QFF.

With QFF2 and QFF3, constant FRAT5 inputs by time frame were used to estimate sell-up for each family. The fundamental assumption with constant FRAT5 values by time frame is that the WTP of passengers in each family does not change significantly over the duration of the booking process. However, it could very well be true that both business and leisure passengers' WTP *does* change as the departure date nears. Thus perhaps a more rigorous method of representing passenger's willingness-to-pay (and consequently their sell-up probability) could make the QFF methods even more effective.

The second suggestion for future research is a direct follow-up to the weaker performance of the QFF3 formulation in competitive RM simulations with overlapping structures. As described in Chapter 5, because the historical data from each family consisted of both business and leisure passengers (more so than in with the non-overlapping structure), the Q-Forecasts by family were less accurate. That is, because of the way the leisure/business bookings are scaled, the Q-Forecasts were too high, which resulted in overly-aggressive closure rates with QFF3. Thus, further investigation on the assumption that historical data from each family is an appropriate way to generate the Q-Forecasts could be conducted. Pursuing this direction could especially be useful for experiments with overlapping fare family structures, where the different passenger types are less segmented.

Overall, the contributions from this thesis are believed to be of practical importance to the airlines, with many legacy carriers transitioning to branded fares. That is, developing appropriate RM methods was vital, given that some industry experts believe that fare family structures may be the key to successful airline pricing in the future.

Bibliography

Barnhart, C., Belobaba, P. P., & Odoni, A. R. (2003). Applications of Operations Research in the Air Transport Industry. *Transportation Science*, *37*(4), 368–391.

Belobaba, P. P. (1987). *Air Travel Demand and Airline Seat Inventory Management*. Ph.D. thesis, Flight Transportation Laboratory, Massachusetts Institute of Technology. Flight Transportation Laboratory.

Belobaba, P. P. (1989). Application of a Probabilistic Decision Model to Airline Seat Inventory Control. *Operations Research*, *37*(2), 183–197.

Belobaba, P. P. (1992a). Optimal vs. Heuristic Methods for Nested Seat Allocation. In *AGIFORS Reservations Control Study Group Meeting*. Brussels, Belgium.

Belobaba, P. P. (1992b). The Revenue Enhancement Potential of Airline Revenue Management Systems. In *Astair Proceedings of Advanced Software Technology in Air Transport*, (pp. 143–164).

Belobaba, P. P. (1998). The Evolution of Airline Yield Management: Fare Class to Origin-Destination Seat Inventory Control. In D. Jenkins (Ed.) *The Handbook of Airline Marketing*, (pp. 285–302). The Aviation Weekly Group of the McGraw-Hill Companies.

Belobaba, P. P. (2002). *Airline Network Revenue Management: Recent Developments and State of the Practice*, (pp. 141–156). McGraw Hill.

Belobaba, P. P. (2010). *Passenger Origin-Destination Simulator (PODS): Summary of Processes and Models*.

Belobaba, P. P., & Hopperstad, C. (2004). Alternative RM Algorithms for Unrestricted Fare Structures. In *AGIFORS Reservations and Yield Management Study Group Meeting*.

Belobaba, P. P., Odoni, A. R., & Barnhart, C. (2009). *The Global Airline Industry*. John Wiley and Sons.

Belobaba, P. P., & Weatherford, L. R. (1996). Comparing Decision Rules that Incorporate Customer Diversion in Perishable Asset Revenue Management Situations,. *Decision Sciences*, *27*(2), 343–363.

Botimer, T. C., & Belobaba, P. P. (1999). Airline Pricing and Fare Product Differentiation: A New Theoretical Framework. *The Journal of the Operational Research Society*, *50*(11), 1085–1097.

Boyd, A. E., & Kallesen, R. (2004). The Science of Revenue Management when Passengers Purchase the Lowest Available Fare. *Journal of Pricing and Revenue Management*, *3*(2), 171–177.

Carrier, E. (2003). *Modeling Airline Passenger Choice: Passenger Preference for Schedule in the Passenger Origin-Destination Simulator (PODS)*. Master's thesis, Massachusetts Institute of Technolgy.

Cizaire, C. (2010). Sell-up and Buy-up Inputs: Impact on QFFs Performance. In *PODS Research Consortium*.

Cleaz-Savoyen, R. L. (2005). *Airline Revenue Management Methods for Less Restricted Fare Structures*. Master's thesis, Massachusetts Institute of Technology.

Cooper, W. L., de Mello, T. H., & Kleywegt, A. J. (2006). Models of the Spiral-Down Effect in Revenue Management. *Operations Research*, *54*(5), 968–987.

de Neufville, R., & Odoni, A. R. (2002). *Airport Systems: Planning, Design, and Management*. McGraw-Hill Professional.

d'Huart, O. (2010). *A Competitive Approach to Airline Revenue Management*. Master's thesis, Massachusetts Institute of Technology.

Fiig, T., Isler, K., Hopperstad, C., & Belobaba, P. P. (2010). Optimization of Mixed Fare Structures: Theory and Applications. *Journal of Pricing and Revenue Management*, *9*(1-2), 152–170.

Fiig, T., Isler, K., Hopperstad, C., & Olsen, S. S. (2012). Forecasting and Optimization of Fare Families. *Journal of Revenue and Pricing Management*, *11*, 322–342.

Gorin, T. O. (2000). *Airline Revenue Management : Sell-up and Forecasting Algorithms*. Master's thesis, Massachusetts Institute of Technology.

Lee, S. (2000). *Modeling Passenger Disutilities in Airline Revenue Management Simulation*. Master's thesis, Massachusetts Institute of Technology.

L'Heureux, E. (1986). A New Twist in Forecasting Short-term Passenger Pickup. In *26th AGIFORS Annual Symposium Proceedings*, (pp. 248–261).

Littlewood, K. (1972). Forecasting and Control of Passenger Bookings. *12th Agifors Symposium*, *4*(2), 95–117.

McGill, J. I., & Ryzin, G. J. V. (1999). Revenue Management: Research Overview and Prospects. *Transportation Science*, *33*(2), 233–256.

Reyes, M. H. (2006). *Hybrid Forecasting for Airline Revenue Management in Semi-Restricted Fare Structures*. Master's thesis, Massachusetts Institute of Technology.

Smith, B. C., Leimkuhler, J. F., & Darrow, R. M. (1992). Yield Management at American Airlines. *Interfaces*, *22*, 8–31.

Smith, B. C., & Penn, W. C. (1988). Analysis of Alternative Origin-Destination Control Strategies. In *AGIFORS Symposium Proceedings*, vol. 28. New Seabury, MA.

Tam, C. F. W. (2008). *Airline Revenue Management Based on Dynamic Programming Incorporating Passenger Sell-Up Behavior*. Master's thesis, Massachusetts Institute of Technology.

Vinod, B., & Moore, K. (2009). Promoting Branded Fare Families and Ancillary Services: Merchandising and its Impacts on the Travel Value Chain. *Journal of Revenue and Pricing Management*, *8*, 174–186.

Weatherford, L. R. (1991). *Perishable Asset Revenue Management in General Business Situations*. Ph.D. thesis, Darden Graduate School of Business Administration, University of Virginia, Charlottesville, VA.

Weatherford, L. R. (1999). Forecast Aggregation and Disaggregation. In *IATA Revenue Management Conference Proceedings*.

Weatherford, L. R., & Polt, S. (2002). Better Unconstraining of Airline Demand Data in Revenue Management Systems for Improved Forecast Accuracy and Greater Revenues. *Journal of Revenue and Pricing Management*, *1*(3), 234–254.

Williamson, E. L. (1992). *Airline Network Seat Inventory Control: Methodologies and Revenue Impacts*. Ph.D. thesis, Massachusetts Institute of Technology.

Zickus, J. S. (1998). *Forecasting for Airline Network Revenue Management : Revenue and Competitive Impacts*. Master's thesis, Massachusetts Institute of Technology.