

Evaluation of Forecasting Techniques for Short-Term Demand of Air Transportation

by

Richard Robert Wickham

S.B., Aeronautics and Astronautics
Massachusetts Institute of Technology, 1993

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, 1995

© Massachusetts Institute of Technology, 1995. All Rights Reserved.

Author
Department of Aeronautics and Astronautics
May, 1995

Certified by
Professor Peter Belobaba
Department of Aeronautics and Astronautics
Thesis Supervisor

Accepted by
Professor Harold Y. Wachman
Chairman, Department Graduate Committee

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

JUL 07 1995

LIBRARIES

ARCHIVES

Evaluation of Forecasting Techniques for Short-Term Demand of Air Transportation

by

Richard Robert Wickham

Submitted to the Department of Aeronautics and Astronautics on
May 15, 1995, in partial fulfillment of the requirements for the
degree of Master of Science in Aeronautics and Astronautics

Abstract

Forecasting is arguably the most critical component of airline management. Essentially, airlines forecast demand to plan the supply of services to respond to that demand. Forecasts of short-term demand facilitate tactical decisions such as pricing and seat inventory control—the allocation of seats among the various booking classes.

In this study, an evaluation was conducted of the relative performance of selected forecasting techniques used to predict short-term demand for air transportation. Short-term in this context is defined as intervals less than eight weeks prior to the date of departure. The selected techniques were representative of current practices in the airline industry including simple time series, linear regression, and booking pickup models. Two types of pickup models were analyzed: the *classical* model and an *advanced* model. The set of models was subjected to the same short-term forecasting environment where the historical data was restricted to ten weekly departures and the forecast horizon limited to seven weeks in the future. Eight scenarios were examined to study the effects of varying the size of the historical data set as well as the length of the forecast horizon. Performance was determined on the basis of the relative accuracy of the forecasts measured through the use of selected metrics.

It will be shown that the booking pickup models consistently outperformed the time series and regression models and the *advanced* pickup model produced the best results. Furthermore, it was discovered that increasing the size of the historical data set beyond seven weekly departures did not have a significant impact on the performance of the various models and in most cases the performance of the models deteriorated as the size of the historical data set was increased.

Thesis Supervisor: Peter Belobaba Ph.D.

Title: Assistant Professor of Aeronautics and Astronautics

Acknowledgments

This thesis represents the culmination of my MIT career. It is a celebration of my experiences over the last six years at this institution and a reflection of the many people who have impacted and shaped my life—to whom I am truly indebted.

I must first thank Professor Peter Belobaba, without whom none of this work would have been possible. Indeed, a thesis without a good advisor is tantamount to a ship without a rudder.

“I am wealthy in my friends”

Shakespeare—*Timon of Athens*. Act II Sc. 2

The circle of people—friends and relatives— whose influence, whether deliberate or inadvertent, has impacted this work is certainly quite substantial. Yet, there are those that are foremost in my mind: “Brother” Bruce McKenzie, for taking the time to savor many cups of coffee with me before our ‘third shifts’; Kim Scholle, who has been a constant source of joy—the keeper of my sanity; My dearest friend, Troy Gatcliffe—although thousands of miles away, his inspiration and support is unwavering; The FTL crew; Tom, Gary, Seth, Carsten, Ray, Eric, and Edmund; and my closest and treasured friends: May, Kathleen, Sridhar, Saf, Mikey, JP, Graham and Sean.

Each of us encounter someone in our lives whose angelic kindness leaves an indelible impression on us. For me that person is Mary Ellen Bushnell. I am truly thankful for all that you have done.

I am also grateful to my dear mother and sister, to whom this work is dedicated, for their continued support and love.

Finally, I am grateful to the Grace of God for the countless blessings I have received.

Alas, as you the reader open the first chapter of this thesis, I would have closed the final chapter in my MIT career.

Enjoy!

Table of Contents

| | |
|---|-----------|
| Introduction | 13 |
| 1.1 Objective of Thesis | 15 |
| 1.2 Organization of Thesis | 16 |
| 1.3 The Airline Booking Process | 17 |
| 1.4 The Economics of The Airline Booking Process | 20 |
| 1.4.1 Demand For Air Transportation | 21 |
| 1.4.2 Supply of Air Transportation | 22 |
| 1.5 The Need for Forecasts | 26 |
| 1.5.1 Forecasting for Revenue Management | 29 |
| 1.5.2 Dependent Variables | 30 |
| Literature Review | 33 |
| 2.1 Categorization of Forecasting Techniques | 33 |
| 2.2 Quantitative Methods | 35 |
| 2.3 Airline Forecasting Literature | 36 |
| 2.3.1 Macro-Level Forecasting | 36 |
| 2.3.2 Passenger Choice Modeling Literature | 37 |
| 2.3.3 Micro-Level Forecasting Literature | 37 |
| Discussion of Selected Models | 43 |
| 3.1 Selection Criteria | 43 |
| 3.2 Selected Model Set | 44 |
| 3.2.1 Time Series | 45 |
| 3.2.2 Regression Models | 47 |
| 3.2.3 Combination (Hybrid) Models | 49 |
| 3.3 Summary of Selected Model Set | 54 |

Experimental Procedure & Data Exploration 55

| | | |
|------------|---|-----------|
| 4.1 | Methodology | 55 |
| 4.2 | Experimental Procedure | 55 |
| 4.2.1 | Reconstruction of Models | 56 |
| 4.2.2 | Data Matrix Generation | 56 |
| 4.2.3 | Forecast Procedure | 57 |
| 4.3 | Data Exploration | 67 |
| 4.3.1 | Market Mix | 67 |
| 4.3.2 | Booking Characteristics | 67 |
| 4.4 | The Short-Term Forecasting Environment | 79 |

Presentation of Results 81

| | | |
|------------|----------------------------------|------------|
| 5.1 | Structure of Presentation | 81 |
| 5.1.1 | Summary of Selected Models | 81 |
| 5.2 | Constrained Scenarios | 82 |
| 5.2.1 | Scenario 1 | 82 |
| 5.2.2 | Scenario 2 | 86 |
| 5.2.3 | Scenario 3 | 90 |
| 5.2.4 | Scenario 4 | 94 |
| 5.3 | Unconstrained Scenarios | 97 |
| 5.3.1 | Scenario 5 | 97 |
| 5.3.2 | Scenario 6 | 101 |
| 5.3.3 | Scenario 7 | 104 |
| 5.3.4 | Scenario 8 | 107 |
| 5.4 | Observations | 108 |
| 5.5 | Summary of Results | 113 |

Conclusions 115

| | | |
|------------|--|------------|
| 6.1 | Research Findings | 115 |
| 6.1.1 | The Short-Term Forecasting Environment | 115 |
| 6.1.2 | The Performance of the Selected Models | 116 |
| 6.2 | Revenue Impact of Forecast Errors | 118 |
| 6.3 | Avenues for Further Study. | 119 |

Bibliography 121

List of Figures

| | | |
|-------------|--|----|
| FIGURE 1.1 | The Seat Inventory Control Process | 14 |
| FIGURE 1.2 | The Airline Booking Process | 18 |
| FIGURE 1.3 | Example Hub & Spoke Network | 26 |
| FIGURE 1.4 | Forecasting Applications and Time-Frames Relative to Flight Departure | 27 |
| FIGURE 1.5 | Y Demand for a Specific Flight in December 1994 | 29 |
| FIGURE 1.6 | The Automated Booking Limit System [8] | 30 |
| FIGURE 2.1 | Alternative Forecasting Techniques [9] | 34 |
| FIGURE 4.1 | Constrained Booking Profiles | 60 |
| FIGURE 4.2 | Unconstrained Booking Profile | 62 |
| FIGURE 4.3 | Sample Profiles for Booking Class B Bookings in a Single Market | 68 |
| FIGURE 4.4 | Average Final Bookings of Class B for Short-Haul Flights | 69 |
| FIGURE 4.5 | Average Final Booking in Class B for Medium-Haul Flights | 70 |
| FIGURE 4.6 | Sample Discount Booking Profile for Medium-Haul Markets | 71 |
| FIGURE 4.7 | Average Final Booking in Class B for Long-Haul Flights | 72 |
| FIGURE 4.8 | Comparison Between Constrained and Unconstrained Bookings for a Single Market | 73 |
| FIGURE 4.9 | Average Final Bookings Across the 13 Week Time Frame for Constrained & Unconstrained Data | 74 |
| FIGURE 4.10 | Sample Profile of Class A Bookings for Short-Haul Markets | 75 |
| FIGURE 4.11 | Average Final Bookings in Class A for Short-Haul Flights | 76 |
| FIGURE 4.12 | Sample Booking Profile for Medium-Haul Market: Class A | 77 |
| FIGURE 4.13 | Average Final Bookings in Class A for Medium-Haul Flights | 77 |
| FIGURE 4.14 | Sample Booking Profile for Long-Haul Market: Class A | 78 |
| FIGURE 4.15 | Average Final Bookings in Class A for Long-Haul Flights | 79 |
| FIGURE 5.1 | Theil's Inequality Coefficient for Scenario 1 | 83 |
| FIGURE 5.2 | MAD for Scenario 1 | 84 |

| | | |
|-------------|---|-----|
| FIGURE 5.3 | MPE for Scenario 1 | 85 |
| FIGURE 5.4 | Theil's Inequality Coefficient For Scenario 2 | 87 |
| FIGURE 5.5 | MPE for Scenario 2 | 88 |
| FIGURE 5.6 | MAD for Scenario 2 | 90 |
| FIGURE 5.7 | Theil's Inequality Coefficient for Scenario 3 | 91 |
| FIGURE 5.8 | MPE for Scenario 3 | 92 |
| FIGURE 5.9 | MAD for Scenario 3 | 92 |
| FIGURE 5.10 | Theil's Inequality Coefficient for Scenario 4 | 94 |
| FIGURE 5.11 | MAD for Scenario 4 | 95 |
| FIGURE 5.12 | MPE for Scenario 4 | 96 |
| FIGURE 5.13 | Theil's Inequality Coefficient for Scenario 5 | 98 |
| FIGURE 5.14 | MAD for Scenario 5 | 99 |
| FIGURE 5.15 | MPE for Scenario 5 | 100 |
| FIGURE 5.16 | Theil's Inequality Coefficient for Scenario 6 | 101 |
| FIGURE 5.17 | MAD for Scenario 6 | 102 |
| FIGURE 5.18 | MPE for Scenario 6 | 103 |
| FIGURE 5.19 | Theil's Inequality Coefficient for Scenario 7 | 104 |
| FIGURE 5.20 | MAD for Scenario 7 | 105 |
| FIGURE 5.21 | MPE for Scenario 7 | 106 |
| FIGURE 5.22 | Theil's Inequality Coefficient for Scenario 8 | 107 |
| FIGURE 5.23 | Comparison of Constrained and Unconstrained Final Bookings for a Sample Market | 112 |
| FIGURE 6.1 | Revenue Impact vs. Forecast Bias [1] | 118 |

List of Tables

| | | |
|------------|---|-----|
| TABLE 1.1 | Example of Virtual Class Structure | 25 |
| TABLE 1.2 | Associated Prices for Network Example | 26 |
| TABLE 3.1 | Booking History Matrix | 45 |
| TABLE 3.2 | Data Subset for Models 1, 2, & 2b | 47 |
| TABLE 3.3 | Data Subset for Regression Model | 49 |
| TABLE 3.4 | Data Subset for Classical Pickup Model | 50 |
| TABLE 3.5 | Data Subset for New Twist Pickup Model | 52 |
| TABLE 3.6 | Pickup Sun-Matrix for Departure in 3 Weeks | 53 |
| TABLE 3.7 | Summary of Selected Models | 54 |
| TABLE 4.1 | Hypothetical Booking History Matrix for Booking Class B for Market AAA-BBB | 57 |
| TABLE 4.2 | Summary of Test Scenarios | 59 |
| TABLE 4.3 | Unconstraining Booking Percentages | 61 |
| TABLE 4.4 | Hypothetical Booking Profile | 62 |
| TABLE 4.5 | Unconstrained Booking Profile | 62 |
| TABLE 4.6 | Division of Data Matrix | 63 |
| TABLE 5.1 | Summary of Selected Models | 81 |
| TABLE 5.2 | Theil Performance Ranking for Scenario 1: 2 Week Forecast Horizon | 84 |
| TABLE 5.3 | Theil Performance Ranking for Scenario 1: 7 Weeks Forecast Horizon | 86 |
| TABLE 5.4 | Theil Performance Ranking for Scenario 2: 4 Weeks of Historical Data | 89 |
| TABLE 5.5 | Theil Performance Ranking for Scenario 2: 10 Weeks of Historical Data | 89 |
| TABLE 5.6 | Theil Performance Ranking for Scenario 3: 2 Week Forecast Horizon | 93 |
| TABLE 5.7 | Theil Performance Ranking for Scenario 3: 7 Week Forecast Horizon | 93 |
| TABLE 5.8 | Theil Performance Ranking for Scenario 4: 4 Weeks of Historical Data | 96 |
| TABLE 5.9 | Theil Performance Ranking for Scenario 4: 10 Weeks of historical Data | 97 |
| TABLE 5.10 | Theil Performance Ranking for Scenario 5: 2 Week Forecast Horizon | 100 |

| | | |
|------------|---|-----|
| TABLE 5.11 | Theil Performance Ranking for Scenario 5: Week 7 | 101 |
| TABLE 5.12 | Theil Performance Ranking for Scenario 6: 4 Weeks of Historical Data | 103 |
| TABLE 5.13 | Theil Performance Ranking for Scenario 6: 10 Weeks of Historical Data | 103 |
| TABLE 5.14 | Theil Performance Ranking for Scenario 7: 2 Week Forecast Horizon | 106 |
| TABLE 5.15 | Theil Performance Ranking for Scenario 7: 7 Week Forecast Horizon | 106 |
| TABLE 5.16 | Theil Performance Ranking for Scenario 8: 4 Weeks of Historical Data | 108 |
| TABLE 5.17 | Theil Performance Ranking for Scenario 8: 10 Weeks of Historical Data | 108 |

Chapter 1

Introduction

Since the advent of deregulation in 1978, the US domestic airline market has evolved into one of the most fiercely competitive industries in the world today. The post deregulation chain of events which led to the industry's current situation stemmed primarily from two freedoms given to US carriers through the deregulation act: the freedom to enter or exit any domestic market and the ability to price to what that market would bear. This ushered in a new era in the US industry: one that is beset with vicious fare wars, substantial excess capacity, and consequent reduced margins. The threat of new entrant low cost carriers has remained real for the traditional majors and the race to match these carriers is on.

The science of *Revenue Management* was developed as a direct response to this new environment and is aimed at offering airlines that implement it a significant competitive advantage. In general, Inventory Control in Revenue Management is the process of saving seats (rooms, cars) for the late-booking passenger (guest, driver) [1]. Specific to air transportation, the objective of Revenue Management is to maximize total passenger revenues through the use of seat inventory control to allocate seats optimally among the various fare classes on a given flight. At the core of Revenue Management is the theory of *differential pricing* which implies identifying different groups of consumers and charging each group a different price for a homogeneous product [2]. This practice allows total cost to be covered by total revenues, whereas marginal cost pricing does not. It is estimated

that the practice of Revenue Management can increase revenues by the order of 5%[3]. Figure 1.1 illustrates the typical seat inventory control process.

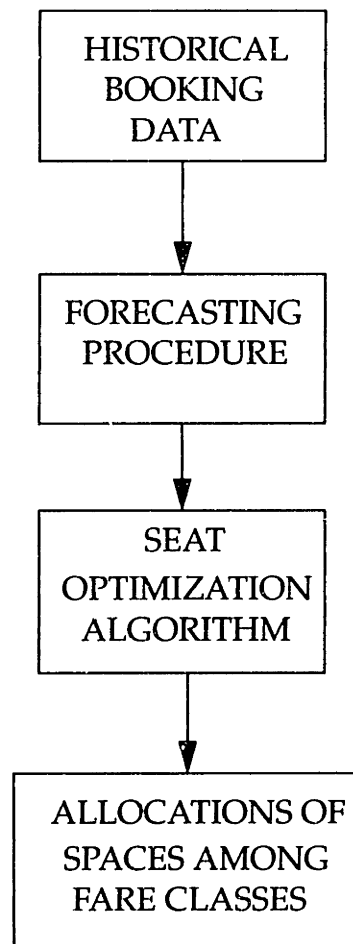


FIGURE 1.1 *The Seat Inventory Control Process*

Central to this process is the need for a forecasting and optimization system. Forecasting and optimization are crucial because they provide the answers to the two questions at the core of the process:

- How many late-booking passengers should be expected? (Forecasting)
-

- How many seats should be saved for these passengers? (Optimization)

Much research has been done on addressing the second question, the optimization of seat allocation among fare classes. However, relatively little attention has been paid to the question of airline reservations forecasting. This is certainly striking since the forecast is a key element of the process and the potential pay-offs of an accurate forecast are substantial, particularly on high demand flights. Studies demonstrate that each 10% improvement in forecast accuracy on high demand flights can potentially translate to a \$10 to \$60 million increase in total annual passenger revenues for a major US airline [4].

1.1 Objective of Thesis

The objective of this thesis is to evaluate the performance of selected forecasting methods used to predict passenger pickup and short term air transportation demand. Passenger pickup is defined as the incremental bookings received during a certain time interval. For example, if on a particular flight the bookings on hand at day 15 and day 10 before departure are 16 and 26 respectively, then the 5-day pickup is 10 passengers.

At a highly aggregate level, forecasting techniques can be classified as either qualitative or quantitative and applied to macro-level, micro-level or passenger choice projections. This study restricts itself to quantitative forecasting techniques used at the micro-level and encompasses regression, time series and hybrid methods (combinations of the two). The chosen methods represent a spectrum of commonly used techniques and will be benchmarked using a fixed data set. Performance is measured on the basis of forecast accuracy as determined by the absolute and percentage errors, inherent bias, as well as factors such as ease of implementation and associated complexity. Based on this benchmarking, a “best practice” protocol is developed to indicate which method is most appropriate or works best under which circumstances.

1.2 Organization of Thesis

The remainder of the thesis is organized in the following manner: In the rest of Chapter 1 the framework for the overall discussion is developed. This entails introducing and defining the concepts and terminology relevant to the airline booking process and Revenue Management. Based on this foundation, the need for forecasting is subsequently distilled.

Chapter 2, Literature Review, presents the salient literature devoted to airline forecasting problems. While most of the literature discusses macro-level forecasting, the primary interest of this thesis lies in micro-level forecasting. Nevertheless, the two levels are described and the general characterization of forecasts delineated.

Chapter 3, provides a detailed presentation of the methods selected for this study. The discussion encompasses the rationale behind choosing these methods as well as certain characteristics such as simplicity, implementation, and data requirements.

Chapter 4, Methodology and Data Exploration, begins by outlining the experimental procedure, including the replication of the various models and the choice of the different forecast scenarios considered. The latter half of the chapter is devoted to the exploration of the data. This includes the overall structure, market mix, booking profiles, seasonality, as well as constrained versus unconstrained demand considerations.

In Chapter 5, Presentation of Results, the discussion centers around forecast errors. The performance of the chosen methods is evaluated based on accuracy, determined by the absolute, mean, and percentage errors. The methods are also scrutinized for inherent biases. Different cases are examined to test the methods over a range of situations within the short term.

Chapter 6 contains the conclusions of this thesis. A “best-practice” protocol based on the benchmarking is created. In addition, this chapter discusses practi-

cal issues surrounding the implementation of an airline reservations forecasting system and the revenue payoffs for increased accuracy.

1.3 The Airline Booking Process

Although airlines essentially market a service to the consumer, the end product is a seat on a given aircraft scheduled to fly from point A to point B at a given time in the future. This product, which can be embellished with meals, drinks, and entertainment, is purchased by the consumer in search of air transportation. The inventory of this product is constrained by the overall fleet size of the respective airlines, as well as the seating capacity of each aircraft, and is therefore fixed. Nevertheless, identical units (seats) can be priced at different levels on the basis of purchase conditions and service amenities, and consequently marketed as distinct service options [2]. Moreover, from the traveller's perspective, the value of the unsold product or empty seat increases as the departure date approaches and reaches a maximum just before departure since most travellers booking at this last minute are willing to pay a premium in order to get a seat on the flight. After departure, the unsold seat can no longer be sold, regardless of the price. In an effort to minimize this risk, airlines have invested heavily in pricing and marketing schemes aimed at stimulating incremental demand.

The airline booking process can be categorized into three phases: the reservation phase, the confirmation phase, and the boarding phase as shown in Figure 1.2 [5]

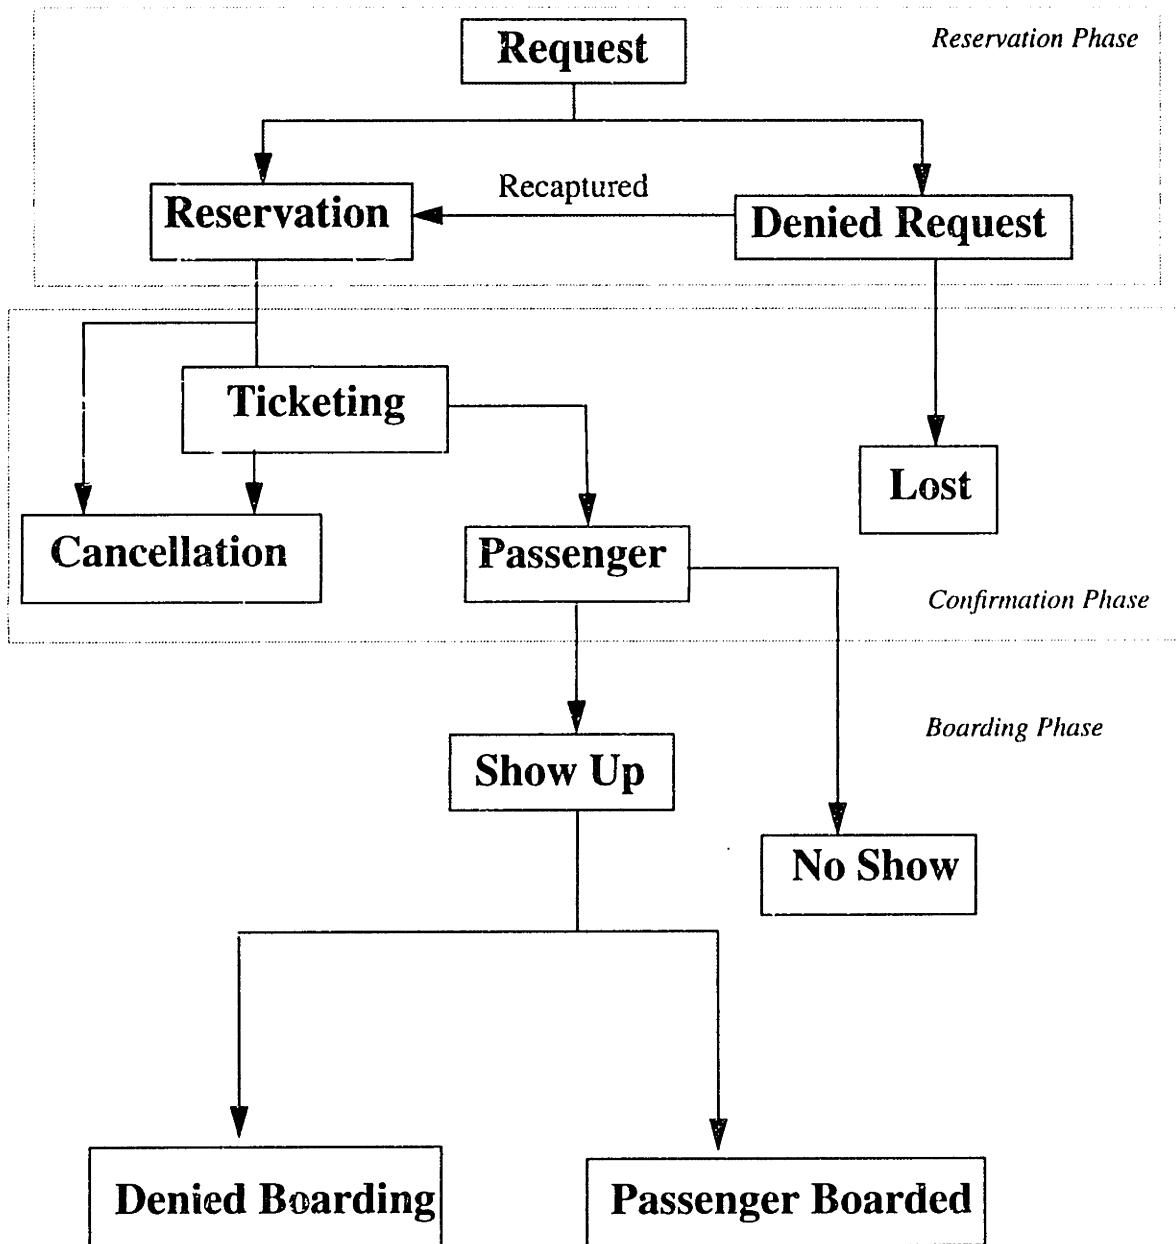


FIGURE 1.2 The Airline Booking Process

The Reservation Phase

In the reservation phase a *request* for air transportation enters the airline's reservations system. A request typically takes the form of a call to a travel agent or

airline reservation agent for air travel services between an origin city and a destination city on a specific date. Once the request is accepted a *reservation* is made and the inventory of available spaces in the given fare class for the given flight is decremented. The distinction between spaces and seats is made since an airline may *overbook* and sell more spaces than physical seats on a given flight, in anticipation that some of the passengers with reservations will not turn up on the day of departure.

If space is not available on the given flight in the given fare class, the request is denied. A *fare class* is a grouping of similar published fares created for the purpose of controlling reservations. On a given flight, each fare class is assigned a certain number of spaces. When the spaces allocated to a certain fare class are filled, the fare class is considered full or closed. Vertical recapture would occur if the traveler whose request was denied is persuaded to accept a different fare class on the same flight. If the traveler is accommodated on another flight in the initially requested fare class, then the airline has made a horizontal recapture [2]. However, if the traveler chooses another airline or decides not to fly at all, the traveler is lost to the airline.

Ticketing occurs when the traveler pays for the service and is given a ticket confirming the itinerary of the trip and the fare class. Depending on the restrictions of the given fare product, ticketing can be done from the time the reservation is made right up to the point of departure.

The Confirmation Phase

The confirmation phase begins immediately after a reservation has been made and continues up to the point of embarkation. During this phase there is the ongoing risk of cancellation, whether explicitly or implicitly. An explicit cancellation occurs if the traveler cancels a reservation or re-books to another fare class or to another flight. An implicit cancellation occurs when the airline cancels a reservation due to the traveler's failure to comply with restrictions. For example,

several discount fares have associated cut-off dates before which they must be ticketed. If a traveler makes a reservation for one of these fares and fails to purchase it before the given date, then the airline would automatically cancel the reservation. In addition, it is well within the airline's prerogative to cancel the reservations of a traveler who does not show up before a certain time on the day of departure.

The Boarding Phase

The boarding phase occurs at the airport on the day of departure. A traveler who shows up with a reservation before the minimum check-in time becomes a *passenger*. If there are sufficient seats, then the passenger is given a seating assignment and allowed to board. However, if more passengers turn up than there are available seats, the flight is considered oversold and some passengers will be denied boarding the aircraft. A denied boarding may be voluntary, where the passenger consents not to board in exchange for some type of compensation. Otherwise, the denied boarding is involuntary, where the airline refuses to accommodate the passenger on the flight. Compensation is then in accordance with the policies of the airline.

1.4 The Economics of The Airline Booking Process

From a microeconomic perspective, the booking process can be viewed as an economic interaction between a consumer (the potential traveler) seeking to maximize utility and a producer (the airline) seeking to maximize profits. It is the consumer who creates the flow through the airline's network by deciding to travel on that specific airline. Collectively, the consumers generate the demand for air transportation. The airline, as a producer of the service, provides a schedule of flights between city pairs with a certain number of available seats in the respective fare classes. Collectively, the world's airlines provide the supply of air transportation.

1.4.1 Demand For Air Transportation

It is important at this point to underscore the distinction between the airline's output, defined as an available seat flown from a point of origin to a point of destination, and the airline product as purchased by the consumer. The demand for air travel is a derived demand, meaning the consumer does not purchase a quantity of available seat miles as if they were a commodity. The value of the purchased air travel is derived from being at a particular place at a particular time.

In keeping with the classical microeconomic theory, travelers will make choices that are most favorable to them. The measure of favorability of a particular alternative is referred to as *utility*. The consumer's main concern therefore, is to maximize utility when requesting air travel. The major factors involved in the decision making process for air transportation include: travel dates, price, service, and restrictions. In general, travelers can be segmented on the basis of the extent to which these factors dominate their choice of service.

Airlines traditionally have identified demand segments under the assumption, supported by empirical evidence, that there exist substantial differences in demand elasticities between business and leisure travelers with little or no cross-elasticities. It has also been revealed that price and service are the critical elements responsible for this segmentation. It is the existence of such segmentation which makes Revenue Management and Differential Pricing possible.

Additional attempts to identify more detailed demand segments have proven to be inherently difficult and complex to the extent that the use of the non-discretionary (business) versus discretionary (leisure) model has become the virtual standard of the industry. This model assumes that discretionary passengers are strongly price sensitive and consequently seek the lowest available fare for any given service. Discretionary passengers are willing to accept restrictions in order to obtain a discount fare. For the non-discretionary passenger, however, quality of service becomes the top priority in choosing a particular flight. As a result, issues such as departure times, frequency of departures, in-flight amenities, and

the ability to make changes to the itinerary are weighted significantly higher than the actual fare. Non-discretionary passengers are willing to pay a premium in order to obtain this desired quality of service.

The demand for air transportation is also characterized by stochastic and seasonal components. Stochastic variation pertains to the inherent volatility of human behavior. The unpredictability associated with the choices travelers make and their responses to certain circumstances adds to the complexity of forecasting. It is this stochastic variation that gives rise to the truism that no forecast of the demand for air transportation can ever be 100 percent accurate.

Seasonal variation occurs quite naturally in the demand for air travel. During certain periods of the year, such as Christmas, summer, or Thanksgiving, surges in demand are typical. In contrast, the period from January to March is well recognized as the “off-season” for air transportation demand. Moreover, at certain times of the year, certain destinations become more desirable. For example, leisure travelers typically target warm climates during the winter months, such as Florida or the Caribbean. Consequently, historical booking data should be “de-seasonalized” before it is used in forecasting, to mitigate the effects of seasonality.

1.4.2 Supply of Air Transportation

The supply of air transportation consists of three major components. The first component comprises a schedule of air services between a set of origins and destinations and is consequently defined over a network of markets. To facilitate the appreciation of this concept and the remainder of the discussion, the fundamental terminology of the air transportation schedule will be defined [6]. A *route map* is a geographical network connecting the cities to be served. A *link* connects two cities in the route map if it is flown non-stop by any aircraft. A *route* is a series of consecutive links flown by an aircraft from origin to final destination with intermediate stops. A *flight* is the passage of an aircraft along a route at some particu-

lar time and is often considered the basic element of supply. A *flight segment* is the portion of a flight over a link. A *route segment* is synonymous with a link.

Fleet assignment represents the second critical component of air transportation supply. Once the route map is defined, the airline must decide which aircraft to assign to each route in its network. In general, the fleet of an airline comprises aircraft of varied physical and performance characteristics, such as seating capacity, cruising speed, maximum range, noise emission level, or minimum takeoff runway requirements. Consequently, factors such as noise abatement restrictions, length of haul, or the length of the runway at the origin airport, would bear on the utilization of a specific aircraft on a specific route. To a large extent, the choice of aircraft determines the operating cost of the flight. From a profit maximizing perspective, however, the seating capacity of the aircraft becomes the prime criterion and literally places an upper constraint on the revenue generating potential of a given flight.

The third major element of air transportation is the method of selling individual seats on the aircraft. An aircraft is typically divided into two or three different cabins, each offering a different level of service and amenities. A typical three cabin configuration comprises a First Class cabin, a Business Class cabin and an Economy (Coach) Class cabin. Recent trends, however, have seen airlines collapse their First and Business Class services into a single and more affordably priced "Business-First" service to stimulate demand in the premium classes. Within the domestic US market, most airlines offer only First Class and Economy Class services on their flights.

Airlines have produced a range of fare products that appeals to the various demand segments in an effort to exploit their revenue-maximizing potential. Moreover, most fare structures are designed to minimize seepage between segments through restrictions on the purchase and use of discounted fares, the most typical of which include advance purchase, round-trip travel, and minimum stay

requirements. These restrictions are often incorporated with capacity controls or limits on the number of seats available to particular fare types.

A published fare for a specific market includes the price, the level of service and any rules, restrictions, or cut-off dates that may apply. For any market, there can be a vast number of published fares. Yet, given the dynamic nature of the airline pricing system, a listing of the published fares being offered in a particular market at any point in time may well become obsolete within 24 hours. This is due to the freedom given to US carriers to make changes to their price structures for domestic origin-destinations markets either instantaneously in their computer reservation systems or overnight through the Air Tariff Publishing Company (ATPCO). Published fares are generally grouped into fare classes for the purpose of controlling bookings in the airline's reservation system. A fare class is designated by a single letter code, such as "Y," "M" or "Q." Although each fare class is assigned to a particular cabin, in most cases there are more fare classes than physical cabins in the aircraft. Consequently, it is quite common for passengers booked in different fare classes to sit in the same cabin (or even next to each other) and receive the same level of in-flight service. For example, one major US carrier uses the codes Y, B, M, H, Q, K, L to designate booking classes within economy class. The Y fare class corresponds to the full Coach fare. The fare classes B through L represent increasingly discounted fares with an increasing number of restrictions.

Recent advances in Revenue Management have introduced the concept of virtual classes where the use of letter codes to represent the various fare classes is replaced by associating numbers with distinct revenue intervals (buckets) as illustrated in Table 1.1

TABLE 1.1 *Example of Virtual Class Structure*

| Revenue Value | Virtual Class ^I |
|----------------|----------------------------|
| \$900 and up | 1 |
| \$850 to \$899 | 2 |
| \$700 to \$849 | 3 |
| \$600 to \$699 | 4 |
| \$450 to \$599 | 5 |
| \$325 to \$449 | 6 |
| \$250 to \$324 | 7 |
| \$200 to \$249 | 8 |
| \$125 to \$199 | 9 |
| \$124 or less | 10 |

I. These are not actual classes but have been created for the purpose of illustration.

This system is network oriented and addresses the problem of comparing the revenue earning potential of individual legs on connecting flights when allocating available seats. For example, consider a network of flights to DCA, connecting through DFW (Figure 1.3) where the number of available seats on the DFW-DCA leg is quite limited.




FIGURE 1.3 *Example Hub & Spoke Network*

TABLE 1.2 *Associated Prices for Network Example*

| Market | 1/2 of RT Fare (Same Letter Class) | Virtual Class | Revenue Bucket |
|--------|---------------------------------------|---------------|----------------|
| LAXDCA | \$265 | 7 | \$250 - \$324 |
| LBBDCA | \$210 | 8 | \$200 - \$249 |
| ACTDCA | \$195 | 9 | \$125 - \$199 |
| DFWDCA | \$190 | 9 | \$125 - \$199 |

Placing the fares into virtual classes reveals that more seats on the DFWDCA flight will be made available to the potential passengers originating in LAX and LBB and fewer to the passengers in ACT and DFW, because the latter fares are in a lower virtual class and therefore contribute less to the system revenues.

1.5 The Need for Forecasts

Forecasting is arguably the most critical area of airline management. Essentially, airlines forecast demand to plan the supply of services to respond to that demand. Short-term forecasts (less than 6 months) facilitate tactical decisions such as

catering, pricing, and seat inventory control. Medium-term traffic forecasts, generally defined as a 6 to 16 month horizon, not only impact the entire operating plan, but also influence the current and upcoming fiscal budgets [7]. Aircraft scheduling decisions, maintenance planning, advertising and sales campaigns, and the opening of new sales offices are among the many decisions which ultimately are dependent on short-term forecasts. However, strategic decisions, such as the creation of new routes or the acquisition of new aircraft, hinge on longer-term forecasts. Figure 1.4 shows the major set of forecast-dependent activities and the associated time frame before departure during which they are applicable.

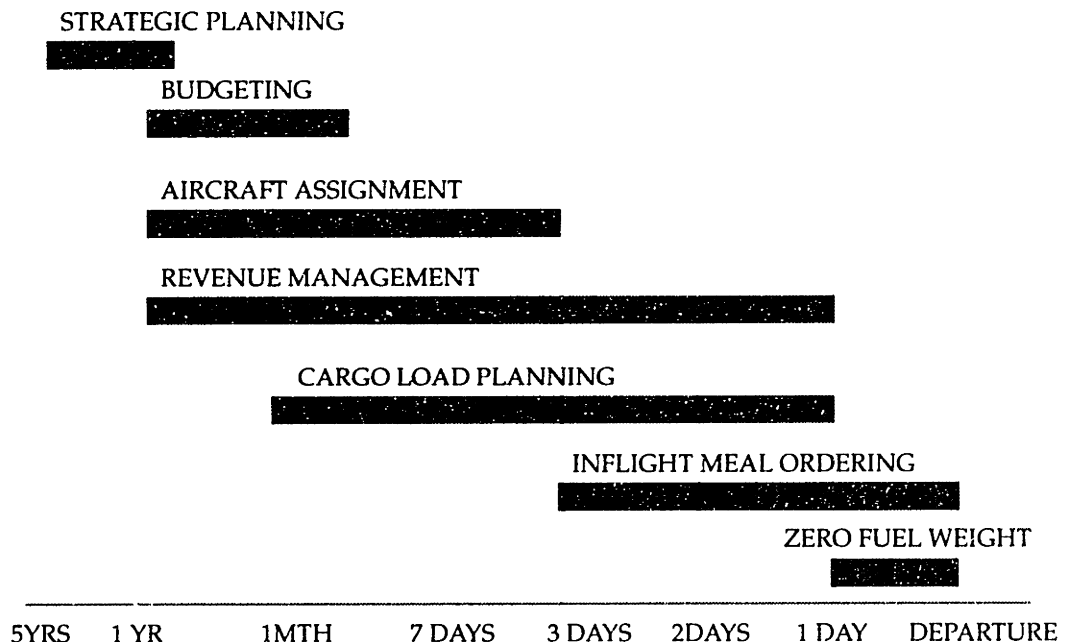


FIGURE 1.4 *Forecasting Applications and Time-Frames Relative to Flight Departure*

There is a mix of circumstances that require forecasting, each of which poses contrasting methodological challenges. For instance, airlines need to forecast traffic growth assuming a continuation of current operating conditions with no drastic changes in fares or in other supply factors. The global growth of passenger and/or freight traffic must be forecasted on a route, group of routes and/or

geographic region basis. From this forecast of total demand, the airline must then predict its own share and corresponding traffic. At the root of such forecasts lies the assumption that traffic growth will continue in the future very much as it has done in the past.

There is also a need to forecast the response of demand to changes in the conditions of supply, such as changes in frequency, capacity, existing fares, or departure times. A significant change in supply conditions may be under consideration by the airline itself or changes may be imposed by competitors. In any event, an airline must be in a position to anticipate the reaction of demand to any such change.

Alternatively, an airline may be faced with the problem of trying to forecast demand on a particular route which is under consideration for new entry. Quite often this may even be a route which has had no previous air service. In any event, the airline has little or no experience nor historical data on which to base its forecasts. Such circumstances make forecasting quite difficult and increase the risk of error. Nevertheless, there are appropriate techniques an airline may utilize in such cases, some of which will be subsequently discussed in Chapter 2.

Lastly, there is the question of segment or disaggregate forecasting. Passenger traffic on a particular flight is composed of distinct market segments related to both travel purpose and service requirements. These segments may be further categorized by point of origin. Studies indicate that each market segment is likely to have different demand elasticities and growth rates. Consequently, it should be possible to achieve more accurate forecasts through aggregating forecasts of each market segment rather than by forecasting total traffic from the start. Some airlines already apply a two-market segment approach to forecasting, namely business and non-business, or devise further segments based on fare classes. In reality, only a handful of airlines currently possess the resources to conduct extensive segmental forecasting. However, the incentive of increased

forecast accuracy has begun to push more airlines to consider disaggregate forecasting.

1.5.1 Forecasting for Revenue Management

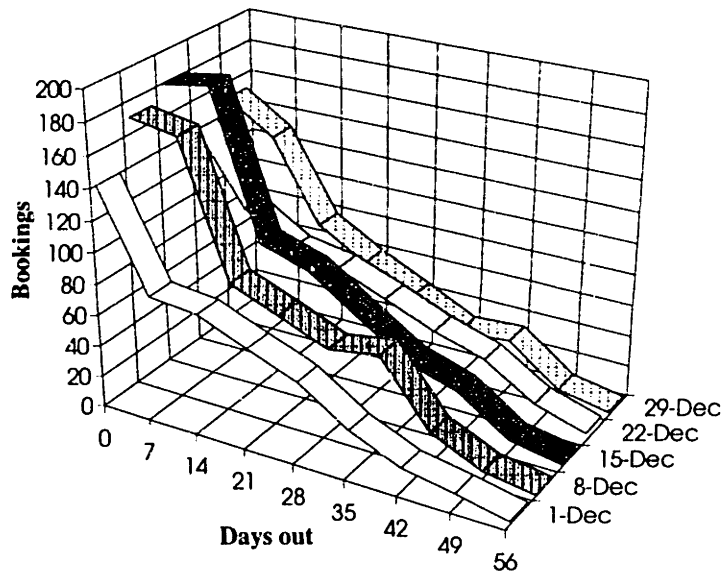


FIGURE 1.5 *Y Class Demand for a Specific Flight in December 1994.*

Forecasting for revenue management is different from traditional forecasting because it involves two time variables. Traditional forecasting uses only one time variable: for example, if the rate of inflation for the last 30 years is known, then it should be possible to predict the rate for the next 5 years.

Revenue Management, however, uses two related time variables: the time of booking (or sale) and the time of consumption. Alternatively, the time of consumption and the days left before consumption can be used. This two dimensional variable space is illustrated in Figure 1.5, which depicts the bookings histories of five departures versus days prior to departure. This figure illustrates the bookings associated with the two time variables of Revenue Management

forecasting: (1) the *consumption date*—the departure date of every Thursday in December, and (2) the days left before this date

The existence of these two dimensions presents one of the many challenges for Revenue Management forecasting.

Given that the airline wants to maximize profit, it requires an accurate forecast of total bookings in each fare class. Figure 1.6, which illustrates the major components in an automated booking limit system, underscores the fact that the forecasting model is central to the entire process

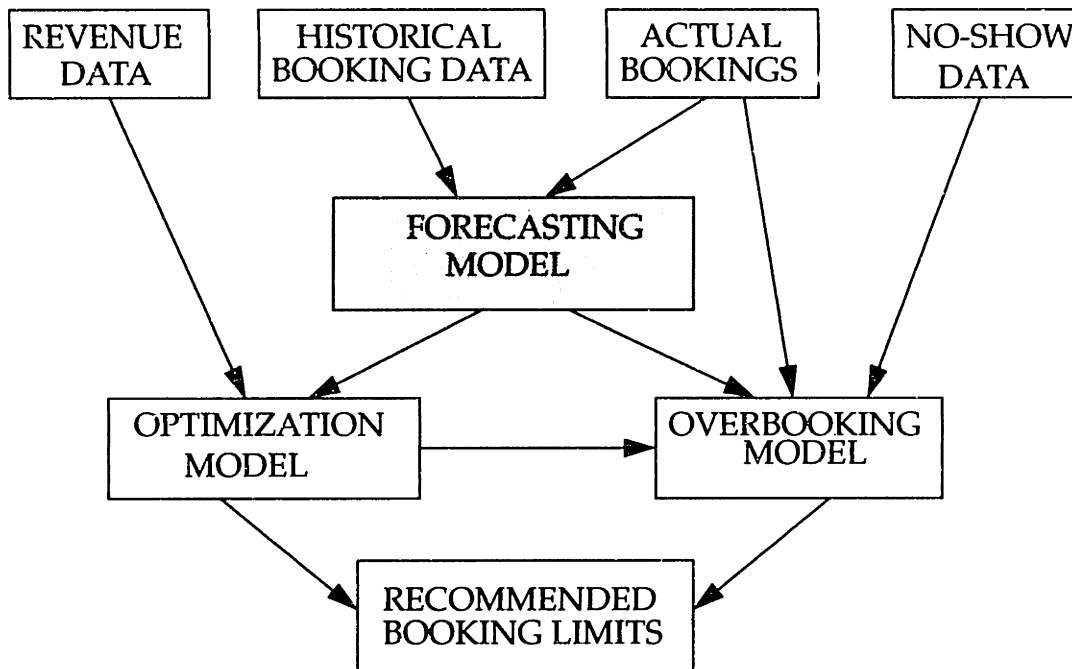


FIGURE 1.6 *The Automated Booking Limit System [8]*

1.5.2 Dependent Variables

Although forecasting demand is often thought of as a major objective of Revenue Management, there are in fact many variables that need to be forecast, including:

-
- **Unconstrained Demand.** Because limits are placed on the number of seats sold in each fare class, the airline only sees constrained data. Unconstrained demand is defined as the number of reservations that would be accepted if restrictions or capacity constraints were not in place. This is certainly difficult to measure and is considered by many to be the “Holy Grail” of Revenue Management forecasting.
 - **Bookings.** Bookings are the actual reservations being held at a particular time. Final bookings refers to the number of reservations being held on the day of departure. This does not, however, indicate the actual load (number of passengers) that would board the aircraft as that number is subject to no-shows and go shows (defined below). The final bookings is a measure of the constrained demand due to the capacity constraints on the aircraft.
 - **Incremental Demand or Pickup.** Some Revenue Management forecasters estimate demand rates, or number of bookings received during certain time intervals. For example, if on a particular flight the bookings on hand at day 15 and day 10 before departure are 16 and 26 respectively, then the 5-day pickup is 10 passengers.
 - **No-Shows.** A no-show is a last minute cancellation by a passenger. In general, a no-show is a traveler with a reservation who fails to show at the airport on the day of departure.
 - **Cancellations.** Cancellations are similar to no-shows, although there is usually time remaining before departure to resell the seat.
 - **Go Shows or Walk-Ups.** These are travelers who show up at the last minute, without reservations, and are willing to purchase a seat.
 - **Sell-Ups and Recaptures.** Sell-ups are travelers who, after having their initial request denied, purchase a seat in a higher fare class. Price elasticity is the underlying consumer characteristic here, the measurement of which poses special problems. Recaptures are rejected requests or reservations
-

whose revenue is not lost to the airline but who purchase a seat on another flight.

This thesis will focus on the forecasting of unconstrained final bookings for a given fare class for a given flight.

Chapter 2

Literature Review

2.1 Categorization of Forecasting Techniques

In general, the forecasting techniques used by airlines can be divided into three broad categories: *qualitative* or judgmental, *quantitative* or scientific, and *decision analysis*, which is a combination of the first two methods (Figure 2.1)[9].

These techniques may be applied at the macro-level, to passenger choice modeling, or micro-level. Examples of macro-level forecasts include projections of total annual domestic traffic and the growth in passenger movements between the US and Europe over the next five years. Passenger choice modeling is the process of predicting an individual passenger's behavior or decision based on socioeconomic factors and the characteristics of alternative options and/or modes for travel. For example, passenger choice modeling can be employed to determine whether an individual would choose rail over air transportation, or choose one airline over another. Micro-level forecasting—the focus of this thesis— pertains to predicting passenger demand at a more detailed or specific level. For example, micro-level forecasting is typically conducted on a flight, date and fare class basis.

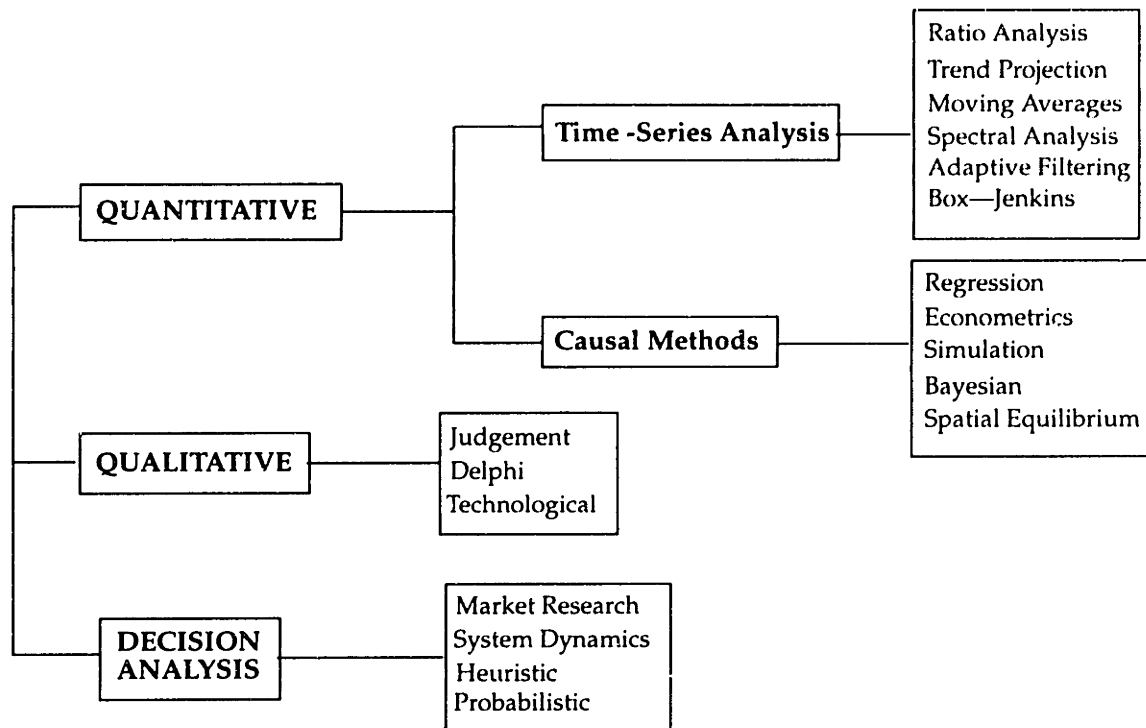


FIGURE 2.1 *Alternative Forecasting Techniques [9]*

In the definition of any forecasting problem at any level, the following three time elements become crucial: the forecasting period, the forecasting horizon, and the forecasting interval[10].

The *forecasting period* is the basic unit of time for which forecasts are made. For example, a forecast may be generated for passenger demand by week, in which case the forecast period is a week. The *forecasting horizon* is the number of periods in the future covered by the forecast. Therefore, if a forecast is required for the next 10 weeks broken down by week, the period is once again one week and the horizon is ten weeks. Sometimes the term lead time is used in place of forecast horizon. Finally, the *forecasting interval* is the frequency with which the new forecasts are generated. Quite often the forecast interval coincides with the forecast period such that the forecasts are revised each period using the most

recent period's demand and other current information as the basis for revision. This would occur, for example, if both the forecasting interval and forecasting period for a particular flight is one week.

In this study, the forecast horizon is restricted to less than 8 weeks and the period varies within this range. The details of this will be discussed in Chapter 4.

2.2 Quantitative Methods

Quantitative forecasting techniques rely heavily on the existence of historical data and, to a large extent, on the continuation of historical trends. This group is divided into two classes: time series analyses and causal methods. Time series analysis tools include methodologies such as ratio analysis, trend projection, moving averages, spectral analysis, adaptive filtering and Box-Jenkins. Detailed discussions on these methods are presented in Montgomery and Johnson [11], Box and Jenkins [12], Brown [13], Jenkins and Watts [14], Anderson [15], and Granger [16]. Causal methods range from regression models to Bayesian analysis.

A time series is a time-ordered sequence of observations of a variable. Time series analysis uses only the time series history of the variable being forecasted in order to predict future values.

Trend projection is the oldest and simplest application of time-series analysis. For example, the demand for wide body aircraft can be estimated as a function of time. Henning [17] further divides this technique into the following three methods:

1. The mean variation method, which derives the forecast from an analysis of various growth rates (e.g. linear, or exponential)
 2. The sliding average method, for which the time-series forecast points can be approximated by an analytical function of just a few neighboring values
-

3. The trend functions method, which draws upon linear, parabolic, logarithmic, or logistic functions to describe the development of the trend.

These methods are based on the premise that what has happened in the past has great relevance to the future. The weakness of this method is that it fails to incorporate the determinants of demand. The impact of changes in the demographic, socioeconomic, and air transportation system variables on air travel is difficult to ascertain.

Nevertheless, time-series analysis is considered especially useful in producing short-term forecasts of monthly, weekly, daily, and hourly variations in demand. Although in the past, the most common methods for dealing with fluctuating patterns have been simple exponential smoothing techniques, significant developments have also been made in techniques such as adaptive filtering, Box-Jenkins methods, and spectral analysis. Examples of applications of these methods in the airline industry can be found in Garvett [18] and Taneja [9].

2.3 Airline Forecasting Literature

2.3.1 Macro-Level Forecasting

At the macro-level of airline forecasting, the principal references are Taneja [9] and Kanafani [19]. In his book, *Airline Traffic Forecasting*, Taneja focuses on regression models for aggregate airline traffic forecasting. He presents statistical methods for macro-issues such as forecasting total airline traffic (on a regional, national, and international scale) and projections of traffic growth. In addition, Taneja argues that causal methods, particularly regression, are the most popular methods of forecasting demand for air transportation. Pure time-series is considered “generally statistical.” This implies that, from a forecasting point of view, methods in this class may answer the “when” question, but do not address the “why” question. Taneja explains that these methods may, for example, be able to predict quite accurately the level of airline passenger demand in 1995, but not

explain why it will be at that particular level. These methods cannot, for example, assess the impact of a reduction in fares, the introduction of new aircraft, an economic recession, or the uncertainties associated with the future labor climate. He contends that such questions can only be answered if the forecaster has specified and calibrated a formal model that shows the influence of all the relevant variables and not just one (i.e. time). This argument is certainly compelling when the forecasting horizon is considerably larger (beyond 10 months). Yet, when the horizon is reduced to a short-term of two months, the probability of drastic variation among the exogenous variables is also reduced. Therefore, the demand characteristics should depend less on these external variables and thus the virtues of causal methods within the short-term are not as clearly defined.

Kanafani addresses in one chapter the issues of aggregate measures of air travel activity such as passenger volume, aircraft operations, and revenue passenger miles. He contends that these measures can be delineated according to trip purpose, origin-destination, length of haul, and type of service (airline, charter, and commuter aviation). The idea of forecasting by fare type is also briefly discussed.

2.3.2 Passenger Choice Modeling Literature

Kanafani [19] offers a brief treatment of passenger choice models in his chapter on demand in air transportation. A categorization of the types of choices which occur in air transportation is developed and includes route choice, airport choice, airline choice, and fare-type choice. A multinomial logic model is presented as a method of estimating passenger choice models. A more general reference to discrete choice modeling in transportation is Ben-Akiva and Lerman [20].

2.3.3 Micro-Level Forecasting Literature

As stated earlier, not much research has been done on micro-level forecasting by flight number, day of week, time of day, or by fare class in the short-term. Little-

wood [21] and Scandinavian Airlines [22] studied some of the basic characteristics of the airline booking process and proposed simple forecasting models for total bookings on a flight. These models are based on computing the mean of historical bookings on previous departures of the same flight. Although Littlewood and Scandinavian Airlines allude to the fact that these models could be used to forecast demand by fare class, the focus is forecasting total demand for the entire cabin on a particular flight, and the emphasis is certainly on simplicity. The Scandinavian Airlines paper also addresses the question of the quantity of historical data necessary to produce accurate forecasts as well as the issue of removal of *outliers* corresponding to unusual, non-recurrent events, such as a promotional sale or the effects of the Gulf war.

The underpinnings of the Scandinavian Airlines paper lie in a study conducted by Duncanson [23] while at Scandinavian Airlines System (SAS). In this study, he reviewed short-term forecasting at SAS and proposed incorporating seasonal analysis and exponential smoothing into the existing models. At that point in time the scope was considered still quite limited and the focus was only applied to passenger traffic with a forecasting horizon of 1 to 3 months. The model was based on historical time series analysis and was directed primarily to relatively stable markets with particular attention paid to European traffic. Duncanson also looked at additive bookings models, as proposed by Littlewood, but did not include cancellations nor day of week effects.

Within academe there are four relevant studies. In his thesis, Sa [24] performed a rudimentary data analysis based on time series models and regression models. Two ARIMA time series models were created for a single fare class on a single flight number. Discouraging results from these models led Sa to subsequently abandon the discussion on time series. His regression model, however, gave more positive results. The dependent variable was bookings to come while the explanatory variables included bookings on-hand, a seasonal index, a day of week index, and a historical average of bookings to come. Nevertheless, Sa did

not test the forecasting ability of the models and therefore the actual predictive abilities of his model remain speculative. In addition, he did not take into consideration the effect of the data being constrained through booking limits.

Brummer et. al.[25] produced the second relevant study, the objective of which was to identify the mean and standard deviation of the true unconstrained log-normal distribution of demand, given a data set with some constrained observations. This study explicitly factors the effect of data constrained by maximum authorized booking limits, although the majority of effort is spent on the derivation of the likelihood function of a censored log-normal distribution and focused only on total bookings on each flight. No attempt is made to study forecasting by class nor is there any attempt to validate the developed model with a different data set.

Research by Ben-Akiva et. al. [26] provides the third relevant study on micro-level forecasting. Three models are proposed to performed flight-specific, class-specific demand forecasting: a regression model for advanced bookings on a given flight, a time series model for historical bookings on previous departures of the same flight number, and a combined model using both advanced bookings and historical bookings data. The preliminary analysis is performed using monthly airline data by flight and fare class. The results indicated that the combined model outperforms both the advance bookings and historical bookings models. Again, even though the results suggested potential for practical application, Ben-Akiva did not have sufficient data to validate the results of the estimated model on future flights. Moreover, the period of the data is monthly, while accurate micro-level forecasting requires data on a weekly if not daily basis. In addition, the effects of constrained demand due to booking limits were not taken into consideration.

Lee devotes his doctoral thesis to developing a comprehensive mathematical framework for the analysis of the airline booking process. His approach uses the work from Rothstein [27] as a basis to develop a complex probabilistic model of

the airline booking process. Unlike Rothstein, however, he considers a stochastic process with interspersed reservations and cancellations, viewed as immigrants and deaths to the population of travelers respectively. The end result is a censored Poisson model for the booking process. A rigorous statistical framework is subsequently developed building on the work of Ben-Akiva et. al. [26]. The effects of booking limits are incorporated using the methodology of Maddala and Schneider [28] to develop a truncated-censored model. Both Maddala and Schneider have done considerable work on the estimation of truncated regression models using normally distributed data. Lee also validates the forecasting ability of his models on actual airline data. The results indicate that the models fit the data well.

Two sources of concise overviews of forecasting for Revenue Management are Belobaba [29] and Curry [1]. In his presentation, entitled *Yield Management Forecasting Made Simple*, to the 4th International IATA Yield Management Conference, Belobaba discussed the importance of forecasting and optimization and outlined the standard industry approaches. In addition, he addressed the issue of revenue benefits derived from accurate forecasts. Curry's technical brief in the Revenue Management Quarterly, *Scorecard*, provided a listing on the tools and techniques as well as related issues such as accuracy factors, revenue impact, and difficulties inherent in the forecasting process.

On the industry front, there are three relevant papers found in the proceedings of the Airline Group of the International Federation of Operational Research Societies (AGIFORS):

1. Harris and Marucci [30] developed a simple regression model in response to Alitalia's product managers request for a method of predicting traffic on their routes in the short term. The model uses two forms of data: (1) 5 historical snapshots of each individual flight taken at 5 different times prior to the day of departure, and (2) a set of data describing the booking situation of all of Alitalia's flights for the next 45 days. The model produces aggregate forecasts for both first

class and economy class as a function of the number of single passenger bookings on hand at a particular point in time as well as the number of groups bookings. In the model, Alitalia's flights are broken down by aircraft type, day of departure, country, continent, and type of flight (i.e. domestic, international, long distance). It was observed that the day of departure did not have significant influence on the model's regression parameters, while all the other factors gave highly significant results. It was also observed that the forecasts for international flights (particularly long-hauls) were far more reliable than those for domestic flights. Nevertheless, this model does not consider the issue of constrained data nor does it address the effects of seasonal variation.

2. Adams and Vodicka [31], while at The Internal Consulting Department at Qantas Airways, reviewed some of the decision making areas in which reliable passenger forecasts are beneficial, namely; operational decisions of cargo capacity planning, in-flight meal ordering, and zero fuel weight estimation. The forecasting horizon for this study was 0 to 7 days prior to departure. Several forecasting models were developed in response to the management's need for various types of information, such as projection reports, threshold curves indicating the variation of the forecasts, and station manager's reports. The emphasis was on simplicity and providing timely solutions. Consequently the models were not very sophisticated and ranged from arithmetic means of segment class load variations exhibited on historical flights to subjective estimates from marketing experts.

3. Of particular interest to this thesis is the paper by Ed L'Heureux [32]. While working for Canadian Pacific Airlines, he developed a "new twist" in forecasting short-term passenger pickup. This new twist builds on the classical pickup model which estimates the pickup for future flight by taking the average of the pickup on previously departed flights. Therefore, on a given day X, using the classical model to forecast the final bookings for a particular flight in the future would require summing the bookings on hand on day X and the estimate of bookings to

come, or pickup, between day X and the day of departure. L'Heureux suggests that the classical method does not exploit the use of the most recent available data and consequently violates the basic maxims in forecasting: "use all of the data" and "give the most weight to the most recent data." This recent data is found in the bookings on flights that have not yet departed. The key to the new twist is to estimate the pickup in smaller increments and sum them to arrive at the pickup for the longer period. In so doing, the data from flights that have not yet departed can easily be incorporated. L'Heureux contends that his new twist model is less influenced by irregular flights such as flights during Christmas time. On the other hand, L'Heureux considers his model to be affected by periods of odd booking activity, such as during fare wars or when a competitor exits a market leaving a surplus of demand. In addition, the new method is said to respond to variations in demand more rapidly, as a direct consequence of using the most recent data. The details of this approach will be discussed in Chapter 3.

Chapter 3

Discussion of Selected Models

As Discussed in Chapter 2, quantitative micro-level forecasting methods can be classified as either time series, regression, or a combination/variation of the two.

Consequently, for the purpose of this comparative study—where the objective is to determine the relative performance of the three classes of forecasting methods for Revenue Management—it is necessary to study at least one model from each of these classes.

3.1 Selection Criteria

In determining which models should be chosen for this study, the following selection criteria were applied:

- **Simplicity:** Based on the industry literature, it is clear that the simplicity of the model is certainly a prime concern for short-term forecasting. Simplicity in this context refers to the level of computational complexity involved in generating a forecast. This simplicity criterion is particularly relevant when considering the time series options that span the gamut of complexity ranging from extrapolation of simple means to the use of auto-regressive moving averages.
-

- **Ease of Application:** This criterion is intertwined with the simplicity measure and pertains to the difficulty associated with reconstructing a particular model. The concern is to avoid overly sophisticated models that require considerable amounts of computing resources. Moreover, with highly sophisticated models the risk of inaccurately reconstructing the model is increased and consequently jeopardizes the validity of any conclusions on performance drawn from this study.
- **Representative of Industry Practice:** Given that this study is fundamentally a benchmarking activity, it is logical therefore, that the focus be on models that are currently being used in the airline industry, as opposed to theoretical models that have not been implemented. After all, such a focus would certainly enhance the value of any conclusions drawn from this study as airlines would immediately be able to recognize the relative performance of their current short-term forecasting methods.
- **Representative of the Spectrum of Complexity:** Although simplicity is a significant driver in the selection of the model set, caution must be taken to ensure that the set does not comprise only the simplest models as this would not be representative of the industry practice.

3.2 Selected Model Set

To facilitate the discussion of the selected models, it is necessary to briefly address the structure of the booking data. Table 3.1 illustrates the generic matrix representation of the booking profile in a hypothetical booking class for a given flight. This particular case displays weekly departures over an 11-week period. Week 0 corresponds to today's date while weeks with negative numbers are historical and those with positive numbers are in the future. For example, week -2 refers to a departure two weeks ago while week 2 represents a departure in two weeks time.

TABLE 3.1 *Booking History Matrix*

| Week | Day0 | Day7 | Day14 | Day21 | Day28 | Day35 | Day42 | Day49 | Day56 |
|------|------|------|-------|-------|-------|-------|-------|-------|-------|
| -5 | 25 | 22 | 10 | 5 | 3 | 3 | 2 | 0 | 0 |
| -4 | 30 | 21 | 15 | 17 | 12 | 7 | 3 | 1 | 0 |
| -3 | 23 | 25 | 14 | 9 | 8 | 5 | 5 | 2 | 1 |
| -2 | 40 | 34 | 30 | 16 | 11 | 6 | 3 | 0 | 0 |
| -1 | 35 | 29 | 20 | 12 | 13 | 8 | 3 | 1 | 0 |
| 0 | 39 | 33 | 30 | 21 | 14 | 6 | 4 | 2 | 1 |
| 1 | — | 28 | 22 | 18 | 10 | 5 | 3 | 0 | 0 |
| 2 | — | — | 18 | 11 | 10 | 7 | 4 | 2 | 1 |
| 3 | — | — | — | 15 | 9 | 8 | 6 | 6 | 2 |
| 4 | — | — | — | — | 11 | 7 | 3 | 2 | 0 |
| 5 | — | — | — | — | — | 9 | 8 | 5 | 2 |

The DayN column displays the bookings on hand N days before departure. For example, Day0 refers to final bookings while Day14 pertains to bookings on hand two weeks before departure. Therefore, on the flight which departed 5 weeks ago, there were 25 final bookings in the sample booking class while there were 10 bookings 14 days before departure.

3.2.1 Time Series

Two basic time series forecasting techniques were selected: simple mean of final bookings and exponential smoothing of final bookings.

Model 1: Simple Mean of Final Bookings

This model generates a forecast on the basis of the average of n historical departures. The forecast of final bookings for a departure t weeks ahead using n historical departures is therefore given by:

$$F\hat{b}kd_t = \frac{1}{n} \sum_{k=-1}^{-n} Fbkd_k \quad 3.1$$

Where $F\hat{b}kd_t$ is the estimated final bookings for the departure on week t and $Fbkd_k$ is actual final bookings for a departure on week k .

Models 2 & 2b: Simple Exponential Smoothing of Final Bookings

This model uses the same basis as the simple mean but applies a smoothing average rather than a pure average. The theory of exponential smoothing implies that the most recent data is weighted heaviest by a smoothing constant α . The forecast for the final bookings for a given period t is given by:

$$F\hat{b}kd_t = \alpha \times Fbkd_t + (1 - \alpha) \times F\hat{b}kd_{t-1} \quad 3.2$$

When the smoothing constant has the value, for example, 0.10, the new estimate places 90% weight on the old estimate and 10% weight on the new observation. In general, the choice of the smoothing constant has an impact on the characteristics of the exponential smoothing. Essentially, the response of the forecast to changes in data is a function of the size of α . The smaller the value of α the slower the response. Larger values of α cause the smoothed value to react quickly—not only to real changes but also to random fluctuations. Typically, when the forecasting period is relatively large, the weights (α) sum to unity [11]. This is not the case however, when the period is small and consequently not all of the data used in the model is captured in the smoothed average. To alleviate this problem, it is possible to either (a) force the weights to sum to unity by creating a *customized* smoothing routine specific to each data case or (b) use a relatively high value for α . Because of the complexity involved in creating a customized smoothing routine, option (b) was employed in this study.

Consequently it was decided to study the performance of exponential smoothing by including two smoothing weights ($\alpha = 0.2$ and $\alpha = 0.4$) to capture the effects

of range of applicable values for α as well as address the issue of having the weights sum to unity.

With both the simple mean and exponential smoothing models, the final bookings of departed flights represent the applicable sample of data from the booking profile matrix (as illustrated by the shaded region in the example matrix below).

TABLE 3.2 *Data Subset for Models 1, 2, & 2b*

| Week | Day0 | Day7 | Day14 | Day21 | Day28 | Day35 | Day42 | Day49 | Day56 |
|------|------|------|-------|-------|-------|-------|-------|-------|-------|
| -5 | 25 | 22 | 10 | 5 | 3 | 3 | 2 | 0 | 0 |
| -4 | 30 | 21 | 15 | 17 | 12 | 7 | 3 | 1 | 0 |
| -3 | 23 | 25 | 14 | 9 | 8 | 5 | 5 | 2 | 1 |
| -2 | 40 | 34 | 30 | 16 | 11 | 6 | 3 | 0 | 0 |
| -1 | 35 | 29 | 20 | 12 | 13 | 8 | 3 | 1 | 0 |
| 0 | 39 | 33 | 30 | 21 | 14 | 6 | 4 | 2 | 1 |
| 1 | — | 28 | 22 | 18 | 10 | 5 | 3 | 0 | 0 |
| 2 | — | — | 18 | 11 | 10 | 7 | 4 | 2 | 1 |
| 3 | — | — | — | 15 | 9 | 8 | 6 | 6 | 2 |
| 4 | — | — | — | — | 11 | 7 | 3 | 2 | 0 |
| 5 | — | — | — | — | — | 9 | 8 | 5 | 2 |

3.2.2 Regression Models

Model 3

The basis of this model lies in determining a linear trend between the final bookings for a departure on week t as a function of the bookings on hand at day $7t$ within the same booking class, as described in the following equation:

$$F\hat{b}kd_t = \beta_0 + \beta \times Bkd_{Day7t} \quad 3.3$$

Where $F\hat{b}kd_t$ is the final bookings for a departure on week t and Bkd_{Day7t} is the bookings on hand at $7t$ days before departure.

A simple least squares regression analysis (with final bookings as the dependent variable and the bookings at day 7t as the explanatory variable) is used to estimate the constants β_0 and β . This study does not address multi-variate regression models.

Model 3b

This model is fundamentally the same as model 3 except that the final bookings for a departure on week t are estimated as a function of the bookings on hand on day 7t in a higher booking class that is representative of a full fare. The motivation for this model arises from the assumption that the bookings in all fare classes are interrelated. Yet, including all the higher booking classes as explanatory variables in a single model is outside the scope of this thesis as it is not intended to study multivariate regression models. As a result, it was decided to utilize a representative higher booking class for a first order analysis of the relationship between the discount and full fare booking classes. The final bookings in booking class Y_x for a departure on week t, is given by the following equation:

$$Fb\hat{k}d_{tY_x} = \beta_0 + \beta \times Bkd_{Day7tY_u} \quad 3.4$$

where Y_u is a higher “full fare” booking class.

The *applicable* data subset for the regression models encompasses a greater fraction of the booking matrix as compared to the time series models. Applicable in this context refers to any data that can be potentially used in the model. For example, given that the largest forecast horizon is 5 weeks ahead using the sample matrix, the applicable data subset contains the bookings from Day35 to Day0 for departed flights (the shaded region in Table 3.3), while the data used to forecast the final bookings for a departure on week 4 is at Day28 and Day 0 only (as shown by the darker shaded columns).

TABLE 3.3 Data Subset for Regression Model

| Week | Day0 | Day7 | Day14 | Day21 | Day28 | Day35 | Day42 | Day49 | Day56 |
|------|------|------|-------|-------|-------|-------|-------|-------|-------|
| -5 | 25 | 22 | 10 | 5 | 3 | 3 | 2 | 0 | 0 |
| -4 | 30 | 21 | 15 | 17 | 12 | 7 | 3 | 1 | 0 |
| -3 | 23 | 25 | 14 | 9 | 8 | 5 | 5 | 2 | 1 |
| -2 | 40 | 34 | 30 | 16 | 11 | 6 | 3 | 0 | 0 |
| -1 | 35 | 29 | 20 | 12 | 13 | 8 | 3 | 1 | 0 |
| 0 | 39 | 33 | 30 | 21 | 14 | 6 | 4 | 2 | 1 |
| 1 | — | 28 | 22 | 18 | 10 | 5 | 3 | 0 | 0 |
| 2 | — | — | 18 | 11 | 10 | 7 | 4 | 2 | 1 |
| 3 | — | — | — | 15 | 9 | 8 | 6 | 6 | 2 |
| 4 | — | — | — | — | 11 | 7 | 3 | 2 | 0 |
| 5 | — | — | — | — | — | 9 | 8 | 5 | 2 |

Data Used To Forecast Final Bookings for a Departure on Week 4

3.2.3 Combination (Hybrid) Models

The set of Hybrid models chosen for this study comprises pickup models only. The generic pickup model implies that the final bookings for a flight departing on week t is a function of the bookings on hand at a particular day X ($X=7t$) as well as the number of booking anticipated to be *picked up* between the given point in time X and the day of departure. This general pickup model can be expressed as:

$$Fb\hat{k}d_t = Bkd_X + PU_{day(X,0)} \quad 3.5$$

where $PU_{day(X,0)}$ is the estimated pickup between day X and the day of departure. This set of generic pickup models can be further subdivided into the following two categories: *Classical Pickup* models and *Advanced Pickup* models [32]. The underpinnings of these two classes are identical to those of the generic

pickup model defined above. However, the distinguishing feature is found in the historical data set utilized by each method.

Models 4, 5 & 5b: Classical Pickup Models

Classical Pickup Models as defined by Duncanson [23] use booking data from departed flights only.

Therefore, in the sample matrix, the applicable data subset, as shown in Table 3.4, comprises all the bookings from the departed flights while the data used to forecast the final bookings on a particular departure on week t comprises the shaded regions of the Day X and Day0 columns. Furthermore, the applicable data sets from the time series and regression models are subsets of the applicable data set for the classical pickup models.

TABLE 3.4 *Data Subset for Classical Pickup Model*

| Week | Day0 | Day7 | Day14 | Day21 | Day28 | Day35 | Day42 | Day49 | Day56 |
|------|------|------|-------|-------|-------|-------|-------|-------|-------|
| -5 | 25 | 22 | 10 | 5 | 3 | 3 | 2 | 0 | 0 |
| -4 | 30 | 21 | 15 | 17 | 12 | 7 | 3 | 1 | 0 |
| -3 | 23 | 25 | 14 | 9 | 8 | 5 | 5 | 2 | 1 |
| -2 | 40 | 34 | 30 | 16 | 11 | 6 | 3 | 0 | 0 |
| -1 | 35 | 29 | 20 | 12 | 13 | 8 | 3 | 1 | 0 |
| 0 | 39 | 33 | 30 | 21 | 14 | 6 | 4 | 2 | 1 |
| 1 | — | 28 | 22 | 18 | 10 | 5 | 3 | 0 | 0 |
| 2 | — | — | 18 | 11 | 10 | 7 | 4 | 2 | 1 |
| 3 | — | — | — | 15 | 9 | 8 | 6 | 6 | 2 |
| 4 | — | — | — | — | 11 | 7 | 3 | 2 | 0 |
| 5 | — | — | — | — | — | 9 | 8 | 5 | 2 |

For example, suppose we wanted to estimate the final bookings in this virtual class for a flight departing in 3 weeks ($t=3$). Using the generic pickup model, the bookings on hand at day X (Day21 = 15) would have to be added to the estimated 3-week pickup from Day21 to Day0. The classical pickup model uses the data from departed flights only to estimate the expected pickup. This is calcu-

lated by subtracting the average bookings on Day21 from the average bookings on Day0 for a given number of historical flights (n). If n is chosen to be 4, then, using a simple average, the 3-week pickup = 20 and the final bookings in the sample booking class for the departure in three weeks = 15 + 20 = 35.

Therefore, using the classical pickup model, the average pickup between day X and the day of departure is given by:

$$\overline{PU}_{day(X,0)} = \overline{Bkb}_{day0_n} - \overline{Bkd}_{dayX_n} \quad 3.6$$

where \overline{Bkb}_{day0_n} is the average final bookings for n departures while \overline{Bkd}_{dayX_n} is the average bookings on day X for n departures.

The specific method used to calculate this average is the distinguishing factor between models 4, 5, and 5b. Model 4 uses a simple average of n departures:

$$\overline{Bkd}_{dayX_n} = \frac{1}{n} \sum_{k=0}^{-n} Bkd_{dayX} \quad 3.7$$

where Bkd_{dayX} is the bookings held on day X for a particular departure n.

Models 5 and 5b employ exponential smoothing (with $\alpha = 0.2$ and $\alpha = 0.4$ respectively) defined as:

$$\overline{Bkd}_{dayX_t} = \alpha Bkd_{dayX_t} + (1 - \alpha) \overline{Bkd}_{dayX_{t-1}} \quad 3.8$$

where Bkd_{dayX_t} is the bookings on day X for a departure on week t.

Model 6, 7, & 7b: Advanced Pickup Model

As discussed in Chapter 2, L'Heureux [32] contended that the Classical Pickup model does exploit the use of all the most recent booking data. L'Heureux argues that this recent data, found in the booking histories of flights which have not yet departed, can add valuable information about the recent booking characteristics

of these particular flights. Consequently, relative to the sample matrix, the applicable data subset for the *advanced pickup* method, as shown in Table 3.5, comprises all of the available booking data.

TABLE 3.5 *Data Subset for Advanced Pickup Model*

| Week | Day0 | Day7 | Day14 | Day21 | Day28 | Day35 | Day42 | Day49 | Day56 |
|------|------|------|-------|-------|-------|-------|-------|-------|-------|
| -5 | 25 | 22 | 10 | 5 | 3 | 3 | 2 | 0 | 0 |
| -4 | 30 | 21 | 15 | 17 | 12 | 7 | 3 | 1 | 0 |
| -3 | 23 | 25 | 14 | 9 | 8 | 5 | 5 | 2 | 1 |
| -2 | 40 | 34 | 30 | 16 | 11 | 6 | 3 | 0 | 0 |
| -1 | 35 | 29 | 20 | 12 | 13 | 8 | 3 | 1 | 0 |
| 0 | 39 | 33 | 30 | 21 | 14 | 6 | 4 | 2 | 1 |
| 1 | — | 28 | 22 | 18 | 10 | 5 | 3 | 0 | 0 |
| 2 | — | — | 18 | 11 | 10 | 7 | 4 | 2 | 1 |
| 3 | — | — | — | 15 | 9 | 8 | 6 | 6 | 2 |
| 4 | — | — | — | — | 11 | 7 | 3 | 2 | 0 |
| 5 | — | — | — | — | — | 9 | 8 | 5 | 2 |

The key to the advanced pickup method involves estimating the aggregate pickup by summing estimates of the pickup over smaller disaggregate intervals. Therefore, before the advanced pickup method can be employed a pickup submatrix must be generated.

Returning to the example applied to the Classical model, the pickup submatrix for the flight departing in 3 weeks is shown in Table 3.6. The pickup between day X and day $X-7$, $PU_{day(X, X-7)}$, for a particular flight is defined as the difference between the bookings on Day X and Day $X-7$:

$$PU_{day(X, X-7)} = Bkd_{day(X-7)} - Bkd_{dayX} \quad 3.9$$

where Bkd_{dayX} is the bookings on day X . Applying the advanced pickup method, with an average of 4 *data flights* to estimate the pickup for each 7-day interval, the pickup in the 3-week period before departure becomes:

$$3\text{-Week Pickup} = \text{Pickup in Week 3} + \text{Pickup in Week 2} + \text{Pickup in Week 1}$$

where the data subset is illustrated by the shaded region in Table 3.6. The term *data flight* is used to indicate that historical data is taken from both departed and non-departed flights.

TABLE 3.6 *Pickup Sub-Matrix for Departure in 3 Weeks*

| Week | Pickup in Wk. 1 before departure | Pickup in Wk. 2 before departure | Pickup in Wk3 before departure |
|------|--|--|--------------------------------------|
| -5 | 3 | 12 | 5 |
| -4 | 9 | 6 | -2 |
| -3 | -2 | 11 | 5 |
| -2 | 6 | 4 | 14 |
| -1 | 6 | 9 | 8 |
| 0 | 6 | 3 | 9 |
| 1 | — | 6 | 4 |
| 2 | — | — | 7 |
| 3 | — | — | — |
| 4 | — | — | — |
| 5 | — | — | — |

If a simple average is used to calculate the average pickup for each interval, then the pickup in week 3 = 7, pickup in week 2 = 5.5, pickup in week 1 = 4, and the 3-Week Pickup = 16.5

Therefore, using the generic pickup model formula, the final bookings in the sample booking class for the departure in 3 weeks is given by the bookings on hand at Day 21 + the estimate of 3-week pickup: $15 + 16.5 = 31.5$.

The difference between models 6, 7, & 7b lies in the method used to estimate the average pickup during the 7-day interval. Model 8 uses a simple mean of the pickup between day X and day X-7 from n departures:

$$\overline{PU}_{day(X, X-7)} = \frac{1}{n} \sum_{t-1}^{t-n} PU_{day(X, X-7)} \quad 3.10$$

while models 9 and 9b use exponential smoothing with $\alpha = 0.2$ and $\alpha = 0.4$ respectively:

$$\overline{PU}_{day(X, X-7)_t} = \alpha PU_{day(X, X-7)_t} + (1 - \alpha) \overline{PU}_{day(X, X-7)_{t-1}} \quad 3.11$$

where $PU_{day(X, X-7)_t}$ is the pickup between day X and day X-7 for a departure on week t.

3.3 Summary of Selected Model Set

As summarized in Table 3.7, the set of models selected for this study comprises 3 time series models, 2 regression models, and 6 pickup models. Models 1 and 3b are designed to serve as the baseline for comparison while the emphasis is deliberately placed on pickup models for this comparative study. Once the initial benchmarking is completed, the set of models will be reduced before conducting the subsequent detailed studies in this thesis. The details of this reduction of the model set will be presented in Chapter 5.

TABLE 3.7 Summary of Selected Models

| Model # | Classification | Method |
|---------|------------------|--|
| 1 | Time Series | Simple mean of final bookings |
| 2 | Time Series | Exponential smoothing of final bkd ($\alpha=0.2$) |
| 2b | Time Series | Exponential smoothing of final bkd ($\alpha=0.4$) |
| 3 | Regression | Fbkd = $f(\text{bkd}_t)$: same booking class |
| 3b | Regression | Fbkd = $f(\text{bkd}_t)$: different booking class |
| 4 | Classical Pickup | Simple mean of total pickup |
| 5 | Classical Pickup | Exponential smoothing of total pickup ($\alpha=0.2$) |
| 5b | Classical Pickup | Exponential smoothing of incremental pickup ($\alpha=0.4$) |
| 6 | Advanced Pickup | Simple mean of incremental pickup |
| 7 | Advanced Pickup | Exponential smoothing of incremental pickup ($\alpha=0.2$) |
| 7b | Advanced Pickup | Exponential smoothing of incremental pickup ($\alpha=0.4$) |

Chapter 4

Experimental Procedure & Data Exploration

4.1 Methodology

The fundamental premise behind the procedure for this study involves constructing a short-term forecasting environment and reviewing the performance of selected forecasting models within this environment. The key component to this methodology becomes, therefore, the application of all the models to the same data set. Constructing this environment involves constraining the dimensions of the data set by placing bounds on the forecasting horizon as well as the size of the historical data set to be utilized. As discussed in Chapter 3, this has a significant impact on the choice of forecasting models for this study. Typically, the horizon for short-term forecasting is confined to within 8 weeks [22]. The further implications of this short-term environment will be addressed in Section 4.4.

4.2 Experimental Procedure

The experimental procedure can be broken down into the following phases:

1. Reconstruction of Models
 2. Data Matrix Generation
 3. Forecast Procedure
 4. Data Reduction & Error Analysis.
-

4.2.1 Reconstruction of Models

The 11 models identified in Chapter 3 were reconstructed using the application *Matlab* on a Unix workstation.

4.2.2 Data Matrix Generation

The data for this study was obtained from a major North American carrier and includes the booking histories in all booking classes for daily flights in 24 markets over an 18 week period (September 1 to December 31, 1994). The booking classes are taken from the set of virtual classes (discussed in Chapter 1) used by this airline for Inventory Control. A more detailed analysis of the characteristics of this booking data will be presented in the latter half of this chapter.

From the complete set of data, the booking histories in two specific booking classes were extracted for weekly departures on a specific day-of-week—chosen randomly to be day 4 (Thursdays). The two virtual booking classes represent a Full fare and a Discount fare class, thereby allowing the analysis to encompass a range of booking activity. These two booking classes will be referred to as classes A & B respectively in the remainder of the discussion.

The data subset was then used to construct the matrices of booking profiles for the 18 weekly day 4 departures by market and by virtual class. Table 4.1 illustrates the generic form of the booking history matrix.

TABLE 4.1 *Hypothetical Booking History Matrix for Booking Class B for Market AAA-BBB*

| Week | Day0 | Day7 | Day14 | Day21 | Day28 | Day35 | Day42 | Day49 | Day56 |
|------|------|------|-------|-------|-------|-------|-------|-------|-------|
| 1 | 25 | 22 | 10 | 5 | 3 | 3 | 2 | 0 | 0 |
| 2 | 30 | 21 | 15 | 17 | 12 | 7 | 3 | 1 | 0 |
| 3 | 23 | 25 | 14 | 9 | 8 | 5 | 5 | 2 | 1 |
| 4 | 40 | 34 | 30 | 16 | 11 | 6 | 3 | 0 | 0 |
| 5 | 35 | 29 | 20 | 12 | 13 | 8 | 3 | 1 | 0 |
| 6 | 39 | 33 | 30 | 21 | 14 | 6 | 4 | 2 | 1 |
| 7 | 45 | 28 | 22 | 18 | 10 | 5 | 3 | 0 | 0 |
| 8 | 50 | 42 | 18 | 11 | 10 | 7 | 4 | 2 | 1 |
| 9 | 33 | 29 | 21 | 15 | 9 | 8 | 6 | 6 | 2 |
| 10 | 46 | 40 | 29 | 22 | 11 | 7 | 3 | 2 | 0 |
| 11 | 49 | 37 | 25 | 17 | 10 | 9 | 8 | 5 | 2 |
| 12 | 25 | 22 | 10 | 5 | 3 | 3 | 2 | 0 | 0 |
| 13 | 30 | 21 | 15 | 17 | 12 | 7 | 3 | 1 | 0 |
| 14 | 23 | 25 | 14 | 9 | 8 | 5 | 5 | 2 | 1 |
| 15 | 40 | 34 | 30 | 16 | 11 | 6 | 3 | 0 | 0 |
| 16 | 35 | 29 | 20 | 12 | 13 | 8 | 3 | 1 | 0 |
| 17 | 39 | 33 | 30 | 21 | 14 | 6 | 4 | 2 | 1 |
| 18 | 45 | 28 | 22 | 18 | 10 | 5 | 3 | 0 | 0 |

4.2.3 Forecast Procedure

Sample Size

A crucial component in the objective of this thesis relies on quantifying the mean forecast errors for the various models. Consequently, in order to estimate these values to some level of statistical significance, it is necessary to utilize an appropriate sample size. This sample size is determined using the following statistical theory [33]:

Assuming a normal distribution of size n with a computed sample mean \bar{x} , the confidence interval for the sample mean at a $100(1-\alpha)$ percent confidence level is given by:

$$\left[\bar{x} - z\left(\frac{\alpha}{2}\right)\left(\frac{\sigma}{\sqrt{n}}\right), \bar{x} + z\left(\frac{\alpha}{2}\right)\left(\frac{\sigma}{\sqrt{n}}\right) \right] \quad 4.1$$

where σ is the standard deviation of the observations about the sample mean and $z(\alpha/2)$ is obtained from the normal distribution tables. Therefore, to be $100(1-\alpha)$ percent confident that the estimate of the sample mean lies within h units of the true value μ , n must be chosen such that

$$h = z\left(\frac{\alpha}{2}\right)\left(\frac{\sigma}{\sqrt{n}}\right) \quad 4.2$$

or equivalently,

$$n = \sigma^2 \frac{\left[z\left(\frac{\alpha}{2}\right) \right]^2}{h^2} \quad 4.3$$

Based on the literature on short-term forecasting, the standard deviation of the forecast errors is in the vicinity of 35%. Consequently, to be 95% confident that the estimated mean forecast error of the various models are within 10% of their true values ($h = 10$, $z(0.025) = 1.96$, $\sigma = 35$), requires a sample size (n) equal to 47.06 observations.

The sample size for this study was therefore set at 48 observations.

Test Scenarios

Table 4.2 provides a summary of the 8 test scenarios conducted in this study:

TABLE 4.2 *Summary of Test Scenarios*

| Scenario | Hst Data Set | Forecast Horizon | Virtual Class ^I | Day of Week | Data Type ^{II} |
|----------|--------------|------------------|----------------------------|-------------|-------------------------|
| 1 | 8 | 2 to 7 | B | 4 | C |
| 2 | 4 to 10 | 4 | B | 4 | C |
| 3 | 8 | 2 to 7 | A | 4 | C |
| 4 | 4 to 10 | 4 | A | 4 | C |
| 5 | 8 | 2 to 7 | B | 4 | U |
| 6 | 4 to 10 | 4 | B | 4 | U |
| 7 | 8 | 2 to 7 | A | 4 | U |
| 8 | 4 to 10 | 4 | A | 4 | U |

I. A = Full Fare, B = Discount Fare

II. C = Constrained, U = Unconstrained

The data in scenarios 1 through 4 is constrained while in scenarios 5 and 6 the data is unconstrained. Scenarios 1, 3, 5, and 7 study the effects of varying the forecast horizon given a fixed historical data set while scenarios 2, 4, 6, and 8 focus on the importance of the size of the historical data set given a fixed forecasting horizon. Eight departures were used in the fixed historical data set (scenarios 1, 3, 5, 7), consistent with the size recommended by Scandinavian Airlines [22]. Given the limits of the values for the forecast horizon and the size of the historical data set, this test matrix was designed to cover the range of possible scenarios without having to conduct each specific case.

Unconstraining

The original data set obtained from the airline includes constrained booking profiles for particular flights, and therefore represents *constrained demand*. Within the data set, in addition to the actual number of bookings held at each incremental checkpoint (Day56, Day49, etc.) the number of available seats is also recorded. Constraining arises when the number of available seats in a particular

booking class on a given day before departure is zero. In such a case, the class may either be *closed* and the final bookings remain at the current level, or the authorization levels may be increased to accommodate additional demand. In either case, however, the data would still be corrupted due to constraining. As illustrated by Figure 4.1, constraining truncates the booking profiles and gives rise to a plateau-like characteristic.

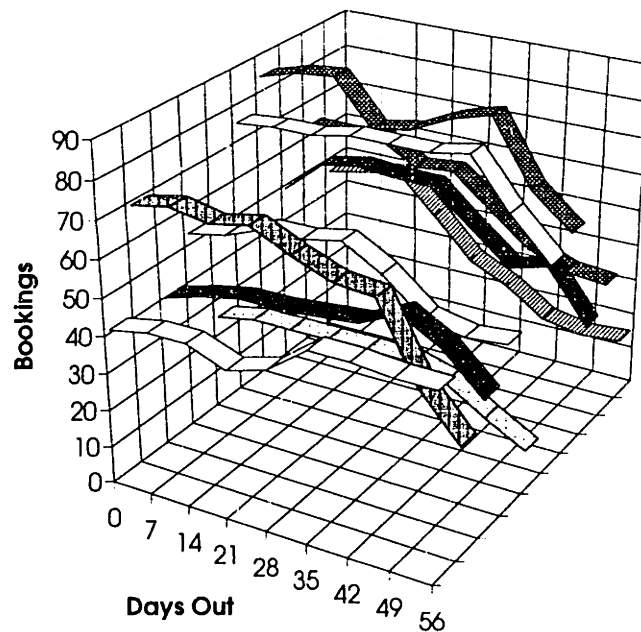


FIGURE 4.1 *Constrained Booking Profiles*

The algorithm for the unconstraining process applied in this thesis is the following:

1. Identify the departures (n) in each market that are not constrained over the entire booking profile for the departure.
2. For these n departures, calculate the average bookings at each interval to produce a single *representative* unconstrained booking profile, given by:

$$B\hat{k}d_t = \frac{1}{n} \sum_{k=1}^n Bkd_{t_{notconstrained}} \quad 4.4$$

Where Bkd_t and $B\hat{k}d_t$ represent the actual and forecasted bookings respectively on day t .

3. Starting at day 56 (as this is the maximum number of days out in the booking matrix), compute the percentage of the bookings at day t relative to the bookings at day $t-7$, given by:

$$\Pi_{t,t-7} = \frac{B\hat{k}d_t}{B\hat{k}d_{t-7}} \quad 4.5$$

4. For a departure, in a given market, with a booking profile constrained at day $t-7$, the unconstrained bookings at day $t-7$ become:

$$Bkd_{t-7_{unconstrained}} = \frac{Bkd_t}{\Pi_{t,t-7}} \quad 4.6$$

5. Repeat step 4 for the bookings on days $x < t-7$, even if they are not constrained, as all data subsequent to the constrained booking at day $t-7$ are considered corrupted.

The incremental unconstraining percentages, $\Pi_{n,n-7}$, computed from the data in the booking class B for this study are shown in Table 4.3.

TABLE 4.3 *Unconstraining Booking Percentages*

| | | | | | | | |
|-------------|--------------|---------------|---------------|---------------|---------------|---------------|---------------|
| $\Pi_{7,0}$ | $\Pi_{14,7}$ | $\Pi_{21,14}$ | $\Pi_{28,21}$ | $\Pi_{35,28}$ | $\Pi_{42,35}$ | $\Pi_{49,42}$ | $\Pi_{56,49}$ |
| 0.934 | 0.897 | 0.849 | 0.826 | 0.848 | 0.823 | 0.815 | 0.841 |

To illustrate the unconstraining process, consider the following example of a booking profile constrained at day 28, as shown in Table 4.3:

TABLE 4.4 *Hypothetical Booking Profile*

| Day0 | Day7 | Day14 | Day21 | Day28 | Day35 | Day42 | Day49 | Day56 |
|------|------|-------|-------|-------|-------|-------|-------|-------|
| 15 | 13 | 11 | 9 | 9 | 8 | 6 | 5 | 2 |

← Corrupted Data →

Using the corresponding percentage $\Pi_{28,35}$ (84.8%) the unconstrained booking for day 28 becomes 9.4. Applying the corresponding percentages to the remaining bookings between day 28 and day 0 produces the unconstrained booking profile shown in Table 4.5.

TABLE 4.5 *Unconstrained Booking Profile*

| Day0 | Day7 | Day14 | Day21 | Day28 | Day35 | Day42 | Day49 | Day56 |
|------|------|-------|-------|-------|-------|-------|-------|-------|
| 16 | 15 | 13.5 | 11.4 | 9.4 | 8 | 6 | 5 | 2 |

Once unconstrained, the previously truncated profiles now behave as growing exponential approaching an asymptotic value (Figure 4.2).

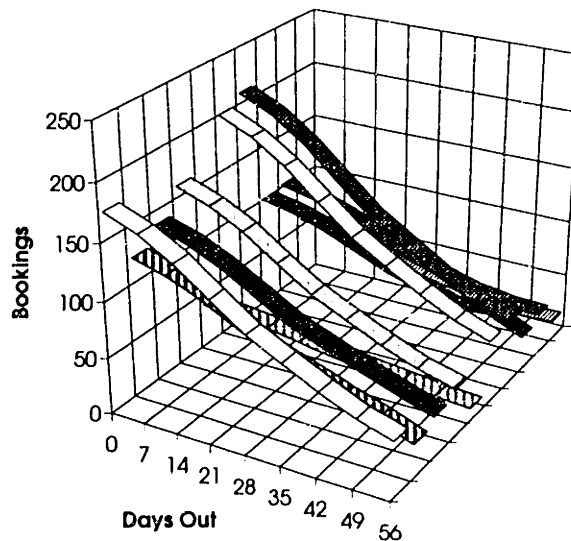


FIGURE 4.2 *Unconstrained Booking Profile*

Data Matrix Manipulation

In order to determine the performance of the various models, the estimated final bookings must be compared to the actual final bookings received for a given departure. Consequently, the need arises for *actual* data to facilitate this comparison. Since the set of booking data comprises historical information only, the actual data is obtained by dividing the individual data matrices into historical and future departures. As illustrated in Table 4.6, this division is accomplished through the use of an artificial *present day* line (week 0).

TABLE 4.6 *Division of Data Matrix*

| Wk. # | Day0 | Day7 | Day14 | Day21 | Day28 | Day35 | Day42 | Day49 | Day56 |
|-------|------|------|-------|-------|-------|-------|-------|-------|-------|
| -11 | 25 | 22 | 10 | 5 | 3 | 3 | 2 | 0 | 0 |
| -10 | 30 | 21 | 15 | 17 | 12 | 7 | 3 | 1 | 0 |
| -9 | 23 | 25 | 14 | 9 | 8 | 5 | 5 | 2 | 1 |
| -8 | 40 | 34 | 30 | 16 | 11 | 6 | 3 | 0 | 0 |
| -7 | 35 | 29 | 20 | 12 | 13 | 8 | 3 | 1 | 0 |
| -6 | 39 | 33 | 30 | 21 | 14 | 6 | 4 | 2 | 1 |
| -5 | 45 | 28 | 22 | 18 | 10 | 5 | 3 | 0 | 0 |
| -4 | 50 | 42 | 18 | 11 | 10 | 7 | 4 | 2 | 1 |
| -3 | 33 | 29 | 21 | 15 | 9 | 8 | 6 | 6 | 2 |
| -2 | 46 | 40 | 29 | 22 | 11 | 7 | 3 | 2 | 0 |
| -1 | 49 | 37 | 25 | 17 | 10 | 9 | 8 | 5 | 2 |
| 0 | 25 | 22 | 10 | 5 | 3 | 3 | 2 | 0 | 0 |
| 1 | 30 | 21 | 15 | 17 | 12 | 7 | 3 | 1 | 0 |
| 2 | 23 | 25 | 14 | 9 | 8 | 5 | 5 | 2 | 1 |
| 3 | 40 | 34 | 30 | 16 | 11 | 6 | 3 | 0 | 0 |
| 4 | 35 | 29 | 20 | 12 | 13 | 8 | 3 | 1 | 0 |
| 5 | 39 | 33 | 30 | 21 | 14 | 6 | 4 | 2 | 1 |
| 6 | 45 | 28 | 22 | 18 | 10 | 5 | 3 | 0 | 0 |

Departures below this line (numbered positively) become future flights— and therefore provide the *actual* data—while departures above this line (numbered negatively) remain as historic departures. Nevertheless, the applicable historical data set comprises the total shaded region shown in the matrix and includes the recent booking data from flights which have not yet departed. The data in the unshaded region of the matrix is not applicable because—given the division of the matrix—it represents bookings which could not have yet been recorded. For example, with a flight scheduled to depart in three weeks time (week 3), it is not possible at this present point in time to know the bookings received on day 7.

As discussed above, the required sample size for this study is 48 observations per forecast. Yet the original data set obtained from the airline comprised 24 markets only. Therefore, once a model is applied to the data matrices for these given markets, the artificial *present day* line, within each matrix, is shifted forward by one week (effectively creating a new historical data set) and the model is then reapplied—generating the 24 additional observations.

Error Analysis

The output from each of the various models consists of the forecasted and actual final bookings for a given scenario as well as the errors or *residuals* defined as the difference between the actual and forecasted values. Scatter plots of the residuals were generated to facilitate the identification of outliers, inherent bias, and covariance. If the models are unbiased, the residuals should be evenly distributed around a mean of zero. Any bias would displace this mean and concentrate the residuals either above or below the zero line. If there is no covariance, there should be no patterns in the scatter plots—the residuals should give the impression that they vary independently within a 2σ horizontal band around the mean.

The following metrics are used to measure performance, where $Fbkd$ and $Fb\hat{k}d$ are the actual and forecasted final bookings respectively generated from n observations:

- **The Mean Absolute Deviation (MAD)**, the average of the absolute values of the forecast errors, is the simplest statistical measure of forecast errors. The MAD is defined mathematically as:

$$MAE = \frac{1}{n} \sum_{k=1}^n abs(Fb\hat{k}d - Fbkd) \quad 4.7$$

The mean absolute deviation is particularly useful when the cost of forecasting errors is proportional to the absolute size of the error.

- **The Mean Percent Error (MPE)** is simply the average of the percentage deviations, defined mathematically as:

$$MPE = \frac{1}{n} \sum_{k=1}^n \frac{(Fb\hat{k}d - Fbkd)}{Fbkd} \times 100 \quad 4.8$$

- **The Mean Absolute Percent Error (MAPE)** is the average of the absolute values of the percentage errors. The mathematical formula for computing the MAPE is:

$$MAPE = \frac{1}{n} \sum_{k=1}^n abs \left[\frac{(Fb\hat{k}d - Fbkd)}{Fbkd} \right] \times 100 \quad 4.9$$

One advantage of this measure is that it is dimensionless. Yet, a particular drawback is that the MAPE is not defined when the actual number of bookings is equal to zero—which is also true for the MPE.

- **The Root Mean Square Error (RMSE)** is the square root of the squared forecasting errors, defined as:

$$RMSE = \sqrt{\frac{1}{n} \sum_{k=1}^n (Fb\hat{k}d - Fbkd)^2} \quad 4.10$$

Where n is the number of observations generated for the particular model. It is important to notice that this measure weighs large forecast errors much more heavily than smaller errors to the extent that it is considered biased against large errors. Nevertheless, it is a valuable measure because of the independence issue.

- **Theil's Inequality Coefficient (U)** [22] is defined as:

$$U^2 = \frac{\frac{1}{n} \sum (Fb\hat{k}d - Fbkd)^2}{\frac{1}{n} \sum Fbkd^2} \quad 4.11$$

Where n is the number of observations generated for the particular model. In this equation the numerator is equal to the mean square error while the denominator is simply the mean square value of the actual final bookings. The numerator captures the actual forecast error whereas the denominator provides a comparison statistic which normalizes the overall coefficient. This metric therefore has the advantage of being dimensionless without the complication of being undefined for zero denominator values (as with the MAPE and MPE). For a perfect forecast U is equal to zero. Consequently, for a particular model, the further away its value of U is from zero the worse the model performs. The Theil's Inequality Coefficient will serve as the primary basis for determining the relative performance of the various models.

The models are ranked on a scenario basis based on the relative values of the calculated metrics. The statistical significance of this ranking is determined using a paired-sample t-test. The details of this method of hypothesis testing can be found in Hogg and Ledolter [34]. For valid hypotheses, the specific differences in the measures are quantified.

4.3 Data Exploration

4.3.1 Market Mix

In choosing the composition of the set of markets for this study, the prime criterion was a sufficient mix of short, medium, and long-hauls to ensure that the data would not be overly biased by any one of the three types of markets. Short-haul markets are defined as distances less than 500 miles, medium-hauls are defined as distances between 500 and 1000 miles, and long-hauls are greater than 1000 miles. The set of 24 domestic markets used in this study comprises 7 short-hauls, 10 medium haul, and 7 long-hauls. In addition, in order to obtain a relatively higher amount of booking activity, the market selection focused on hub movements, where 20 of the 24 markets are flights to and from the major hubs of the carrier.

4.3.2 Booking Characteristics

The set of booking data spans the period from September 1, 1994 to December 31, 1994. At an aggregate level, there is a considerable amount of variation in the characteristics of the booking profiles of the three types of markets over the entire interval. This variation is attributed to seasonality, particularly in light of the holiday periods—Thanksgiving and Christmas—giving rise to significant undulations in demand. The weeks corresponding to these two events are indeed outliers and were consequently removed from the data set. Nevertheless, because the data set does not span an entire year, it is not possible to construct seasonal indices for the specific months within the data set. As a result, it was decided to

neglect the effects of seasonality during this study especially as the data set spans a relatively short time frame displaying the typical seasonal behavior for this time of year (autumn).

At a more disaggregate level, the markets can be separated by length of haul as well as booking class and then compared on a constrained versus unconstrained basis.

Booking Class B (Constrained)

Within the discount virtual class, there are 19 of the 24 markets with at least one constrained flight. 20% of the total number of flights in the data set are constrained.




FIGURE 4.3 *Sample Profiles for Booking Class B Bookings in a Single Market*

The effects of constraining can be seen by the truncated characteristic of the booking profiles, as illustrated in Figure 4.3, where the booking levels appear to saturate before the day of departure.

Booking Class B: Short-Haul

The booking profiles of the discount booking class for the short-haul markets display substantial variability over the 18 week time frame. Overall, the majority of the bookings occur closer to the day of departure, yet, on average, the final bookings are relatively low and seldom go above 25 passengers. The average final bookings for the set of short-haul markets is 8.30 passengers with a standard deviation of 3.6—a 38% variation (Figure 4.4). No significant trends are observed in the behavior of the final bookings.




FIGURE 4.4 *Average Final Bookings of Class B for Short-Haul Flights*

At the origin-destination level, these short-haul routes are primarily business markets—beyond which they serve as hub feeds for connecting leisure travel. Combined, these two characteristics result in the low booking levels. Because the demand for this fare class in these markets is not relatively high the majority of flights remain unconstrained. In fact, the booking profiles indicate that several of the early bookings are frequently lost before the day of departure. When constraining does occur however, it predominantly affects flights within the Christmas holiday period (weeks 15 to 18).

Booking Class B: Medium-Haul

The booking profiles of the medium-haul markets display a more consistent behavior across the time frame with an average final bookings of 25.7 passengers as shown in Figure 4.5.




FIGURE 4.5 *Average Final Booking in Class B for Medium-Haul Flights*

Nevertheless, the standard deviation of 3.6 indicates that the variability of the bookings is comparable to that of the short-haul markets. The majority of the booking profiles appear truncated revealing the presence of booking constraints. Compared to the short-haul markets, the medium-haul booking profiles display greater slopes suggesting the pickup occurs more rapidly and over a shorter period of time. The booking profiles also show the effects of cancellations—dips in the upward sloping characteristic—yet these occur closer to the day of departure Figure 4.6.




FIGURE 4.6 *Sample Discount Booking Profile for Medium-Haul Markets*

The medium-haul markets comprise several leisure destinations which accounts for the overall higher booking levels. Furthermore, it can be argued that the dollar range associated with this virtual class represents a greater discount off the medium-haul full fares as compared to the short-haul full fares. Consequently, given the elastic behavior of leisure travellers, it is anticipated that the relatively greater discount would attract higher demand.

Booking Class B: Long-Haul

The booking levels of the long-haul markets exhibit the most variability over the time frame (Figure 4.7).




FIGURE 4.7 *Average Final Booking in Class B for Long-Haul Flights*

The average final bookings is 20.9 passengers—greater than the short-hauls yet smaller than the medium-hauls. The standard deviation (5.42) however, exceeds both the medium and short-haul markets by over 50%. Seasonality does not appear to be a main contributing factor in this case as the undulations in the bookings across the 18 weeks behave randomly market to market. Furthermore, there does not appear to be any point in the booking period where the majority of bookings consistently occur—some flights receive early pickup, while with others the majority of the pickup occurs closer to the date of departure. The majority of the booking profiles are constrained—as evidenced by the truncated appearance—and show the effects of cancellations prior to departure.

By nature, these long-haul routes represent a mix of business and leisure markets giving rise to the inconsistent booking patterns.

Booking Class B (Unconstrained)

The data from departures constrained by closed booking limits are unconstrained using the algorithm described above (Section 4.2.3). The unconstraining process removes the truncated appearance of the booking profiles and allows the profiles to take on the asymptotic exponential behavior (as illustrated in Figure 4.2) indicative of the *true* demand. As discussed above, the effects of constraining occur consistently towards the end of the 18 week time frame—during the holiday period—for the majority of the markets. Consequently, it is towards this latter part of the time frame where the differences between the unconstrained and constrained final bookings are most pronounced, as illustrated in Figure 4.8.




FIGURE 4.8 Comparison Between Constrained and Unconstrained Bookings for a Single Market

Notwithstanding, at an aggregate level, the average final bookings of the 24 markets remain fairly unchanged by the Unconstraining process (as shown in Figure 4.9), despite the significant changes in the individual markets.




FIGURE 4.9 *Average Final Bookings Across the 18 Week Time Frame for Constrained & Unconstrained Data*

Booking Class A (Constrained)

Overall, only 8 of the 24 markets contain flights with constrained bookings in this full fare booking class. Five percent of the total number of flights in the data set contain constrained booking profiles. This lack of constraining is consistent with business passenger bookings where capacity constraints are seldom effected due to the relatively low volumes of full fare bookings received. In general, the booking levels in the three market types are relatively low and seldom exceed single digit values, in contrast to the associated authorization levels which are typically in the vicinity of 100 bookings. The overall form of the booking profiles in this class is exemplified in Figure 4.10.




FIGURE 4.10 *Sample Profiles of Class A Bookings for Short-Haul Markets.*

Compared to the virtual B class bookings, virtual A booking profiles display more *zero space*—regions where the booking levels are zero. Moreover, because the pickup is relatively small and occurs over a short range, the profiles exhibit a step function characteristic.

Booking Class A: Short-Haul

The majority of the bookings in this fare class are zero over the range of *days out* before departure across the 18 week time frame. On the occasions when non-zero bookings do occur, they are received relatively close to the day of departure and the incremental pickup is marginal. This gives rise to a step function characteristic of the booking profiles for most of the flights in the short-haul markets, although there is no apparent commonality in the appearance of the booking pro-

files from market to market. The average final bookings across the time frame is less than one passenger (0.73) with a deviation of 0.5.

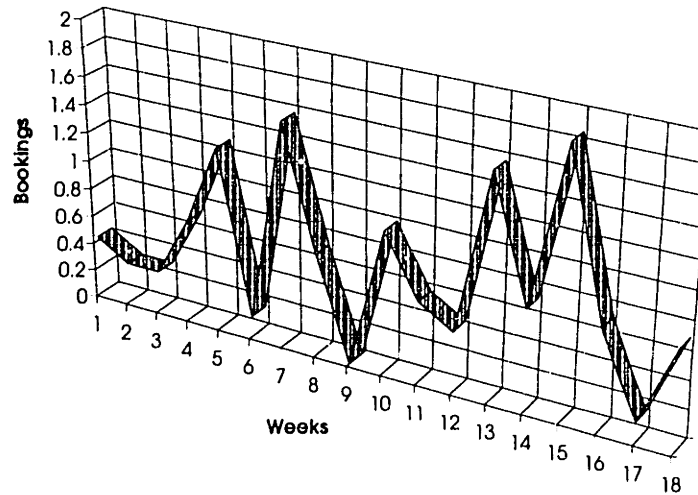


FIGURE 4.11 *Average Final Bookings in Class A for Short-Haul Flights*

Considering the average size of the bookings, this deviation is significant (as shown in Figure 4.11), giving rise to an oscillatory behavior. The fact that these markets are predominantly business destinations and hub feeders accounts for the sporadic booking behavior in the equivalent full fare class.

Booking Class A: Medium-Haul

Once again the booking profiles display considerable amounts of zero bookings over the range of *days out* before departure across the time frame (Figure 12). However, the booking profiles also show a mix of characteristic behavior: some behave as step functions while others display exponential behavior with relatively late pickup.




FIGURE 4.12 *Sample Booking Profiles for Medium-Haul Market: Class A*




FIGURE 4.13 *Average Final Bookings in Class A for Medium-Haul Flights*

As illustrated in Figure 4.13, the average final bookings are relatively high (8.8) and the variation is considerable (3.1). Again, the phenomena discussed in the discount class above is at play: Because the chosen full fare virtual class is the same for the three market types, in certain cases, within the medium-haul markets, it may correspond to the cost of a true full fare while in others it may represent some discount of the full fare value. What is missing from this data is an appreciation of the pricing structure of the individual markets which would allow the dollar value associated with this virtual class to be placed in perspective. Seasonality is not as significant as the variation in the booking profiles appears quite random across the 18 week time frame.

Booking Class A: Long-Haul




FIGURE 4.14 *Sample Booking Profile for Long-Haul Market: Class A*

The booking profiles for the long-haul markets also display the step function characteristic seen in the short and medium-hauls (Figure 4.14). This behavior is

consistent with the relatively low booking levels combined with the low variation. Once bookings are received, the levels remain almost at steady state for the remaining booking duration before departure.




FIGURE 4.15 *Average Final Bookings in Class A for Long-Haul Flights*

The average final bookings over the time frame is approximately one passenger (0.98) with a variation comparable to the short-haul markets (0.53).

Booking Class A (Unconstrained)

Because the majority of the data in this class were originally not constrained, the unconstraining process did not have a significant effect on the booking profiles. The step function characteristic is still evident although the magnitudes of the steps may have increased in certain cases.

4.4 The Short-Term Forecasting Environment

Short-term forecasting can be likened to any fine tuning process involving some type of feedback. Because the parameters are relatively small, the feedback required to achieve the desired accuracies appears more precise. In the case of

this study the parameters are the size of the historical data set, the forecast horizon, and the booking levels while the amount of historical data represents the feedback to the process. In this light, the constraints imposed on the parameters to define the environment can be appreciated.

The historical data set is confined to a maximum of 10 weekly departures. Beyond this point, the additional historical data adds little value in capturing the recent booking characteristics. The 18 week time frame in this study represents a section of a time series of booking data. This interval is sufficiently small that the local variation could be considered independent of the rest of the series. Moreover, as seen in the discussion on the booking profiles, this variation is highly volatile. Under these circumstances, therefore, it is the recent data that would provide the most information on the localized booking behavior from which incremental forecast can be extrapolated.

The forecast horizon is kept to within 8 weeks as this is consistent with the time frame for tactical decisions for Revenue Management.

Based on the exploration of the booking data the following observations have been made:

- The variation of the bookings over the entire short-term environment is truly volatile.
 - Although there is certainly some oscillatory behavior in average final booking profiles, the underlying trends are not apparent.
 - Due to the relatively small time frame, seasonality does not have a significant impact.
 - Unconstraining has a more visible impact on the discount bookings than on the full fare bookings.
-

Chapter 5

Presentation of Results

5.1 Structure of Presentation

As discussed in Chapter 3, eight scenarios are examined in this study, where scenarios 1 to 4 utilize constrained data, while in scenarios 5 to 8, the data is unconstrained. The discussion of the results, presented in the order of the scenarios, will begin with the constrained cases and proceed to the unconstrained cases. The comparison of results will center around Theil's Inequality Coefficient (U) (section), the MAD and the MPE, where U will serve as the primary measure in the evaluation of the relative performance of the models. The RMSE is captured in the definition of U and consequently there is no need for a separate discussion.

5.1.1 Summary of Selected Models

To facilitate the presentation of results, a summary of the set of models used in the various scenarios is given Table 5.1

TABLE 5.1 *Summary of Selected Models*

| Model # | Classification | Method |
|---------|----------------|---|
| 1 | Time Series | Simple mean of final bookings |
| 2 | Time Series | Exponential smoothing of final bkd ($\alpha=0.2$) |
| 2b | Time Series | Exponential smoothing of final bkd ($\alpha=0.4$) |
| 3 | Regression | Fbkd = $f(\text{bkd}_t)$: same booking class |

TABLE 5.1 *Summary of Selected Models*

| Model # | Classification | Method |
|---------|------------------|--|
| 3b | Regression | Fbkd = $f(\text{bkd}_t)$: different booking class |
| 4 | Classical Pickup | Simple mean of total pickup |
| 5 | Classical Pickup | Exponential smoothing of total pickup ($\alpha=0.2$) |
| 5b | Classical Pickup | Exponential smoothing of total pickup ($\alpha=0.4$) |
| 6 | Advanced Pickup | Simple mean of incremental pickup |
| 7 | Advanced Pickup | Exponential smoothing of incremental pickup ($\alpha=0.2$) |
| 7b | Advanced Pickup | Exponential smoothing of incremental pickup ($\alpha=0.4$) |

5.2 Constrained Scenarios

The constrained scenarios are scenarios 1 to 4. Scenarios 1 and 2 pertain to the discount booking class B while scenarios 3 and 4 pertain to the higher booking class A.

5.2.1 Scenario 1

In this scenario, the historical data set (HDS) is fixed at 8 weeks and the forecast horizon (FH) varied between 2 and 7 weeks for the discount booking class B.

Figure 5.1 shows the Theil's Inequality Coefficient (U) for the 11 models over the range of forecast horizons. The overall characteristic of the plot suggests that the U value increases with the forecast horizon indicating that the performance of the models decreases as the forecast horizon increases. A closer look, however, reveals fluctuations where not all of the inter-horizon slopes are positive. In fact, only the U value for Model 5 displays positive growth over the entire range of forecast horizons.

Most of the models appear to be divided between two distinct bands of similar behavior. The first band, comprising the pickup models (4, 5, 5b, 6, 7, 7b), is found on the lower section of the plot where the U values vary between 0.2 and 0.45. The characteristic behavior of the models within this band is a gentle upward slope. The *spread* of the U values among the models is initially quite




FIGURE 5.1 *Theil's Inequality Coefficient for Scenario i*

small but diverges significantly beyond 4 weeks out. Spread in this context is defined as the range of the variation across the models for a particular measure. The second band, comprising the models 1, 2, and 2b, is situated at the top section of the plot and varies between 0.35 and 0.50. The behavior of the models within this band is almost sinusoidal where the spread of the U values diverges between 2 and 4 weeks out, after which it quickly converges and remains relatively tight beyond 5 weeks out. The gap separating the two bands indicate the average differential in performance between the times series and pickup models—25%.

Models 2 and 2b reside for the most part at the extremes if not outside the two bands. Model 3 departs from the first band after 3 weeks out and enters the second after 5 weeks out—although the behavior is not totally consistent with the other models in the second band. Model 3b deviates from the second band when

the forecast horizon exceeds 4 weeks and continues to behave inversely to the other models in this band for the remainder weeks out.

This distinct band behavior is also seen in the MAD (Figure 5.2) where the characteristics are almost identical to the Theil's Inequality coefficient. The lower band of pickup models maintain a MAD of approximately 4 for horizons out to 6 weeks, after which the MAD increases by almost a factor of two.

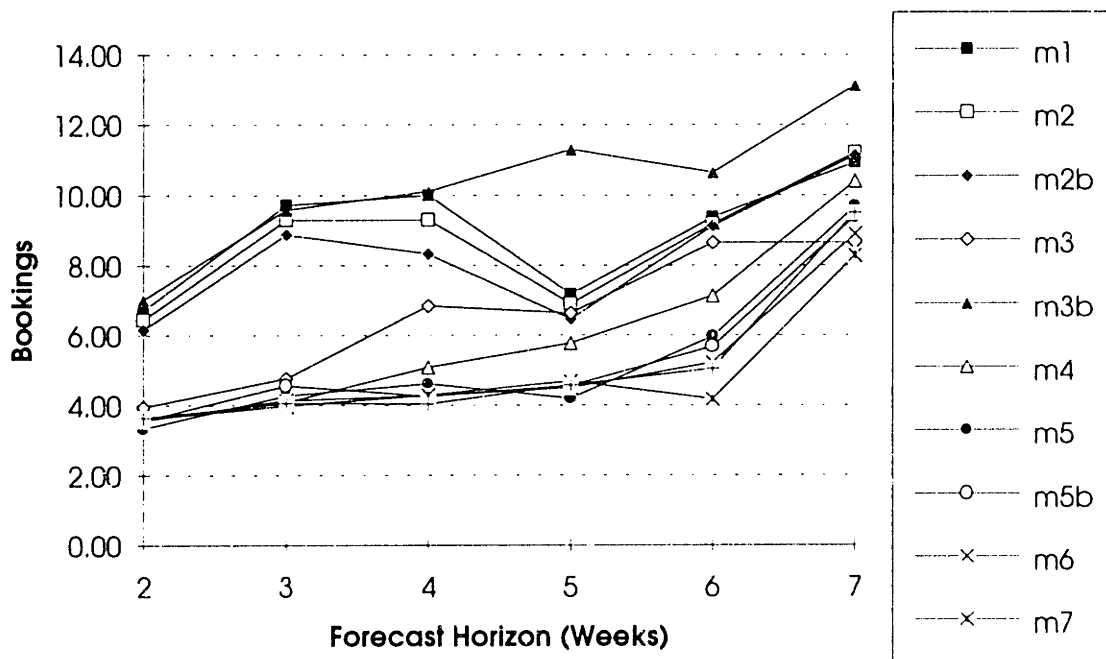


FIGURE 5.2 MAD for Scenario 1

A look at the individual weeks reveals the relative performance of the models. Table 5.2 gives a snapshot of performance metrics for the 2 week horizon where the 11 models are ranked in order of increasing U.

TABLE 5.2 Theil Performance Ranking for Scenario 1: 2 Week Forecast Horizon

| Model # | MPE | RMSE | MAD | MAPE | Theil |
|---------|--------|-------|-------|-------|-------|
| 5 | -0.090 | 5.551 | 3.333 | 0.229 | 0.222 |
| 6 | -0.034 | 6.148 | 3.573 | 0.227 | 0.246 |
| 5b | -0.081 | 6.227 | 3.555 | 0.247 | 0.249 |

TABLE 5.2 *Theil Performance Ranking for Scenario 1: 2 Week Forecast Horizon*

| Model # | MPE | RMSE | MAD | MAPE | Theil |
|---------|--------|-------|-------|-------|-------|
| 7b | -0.057 | 6.244 | 3.643 | 0.273 | 0.250 |
| 7 | -0.049 | 6.286 | 3.605 | 0.263 | 0.252 |
| 4 | -0.028 | 6.079 | 3.583 | 0.269 | 0.278 |
| 3 | -0.013 | 7.254 | 3.942 | 0.250 | 0.291 |
| 2 | 0.144 | 8.694 | 6.448 | 0.496 | 0.348 |
| 2b | 0.056 | 8.762 | 6.165 | 0.469 | 0.351 |
| 1 | 0.211 | 9.039 | 6.716 | 0.523 | 0.362 |
| 3b | 0.242 | 9.522 | 7.020 | 0.553 | 0.381 |

Pickup model 5 is the top performer with respect to the Theil's coefficient and the MAD, while the regression model 3b performs the worst. With respect to the MPE, however, (Figure 5.3), model 3 has the lowest error. In addition, models 1, 2,2b, all show a positive bias, while the biases of the remaining models are negative.

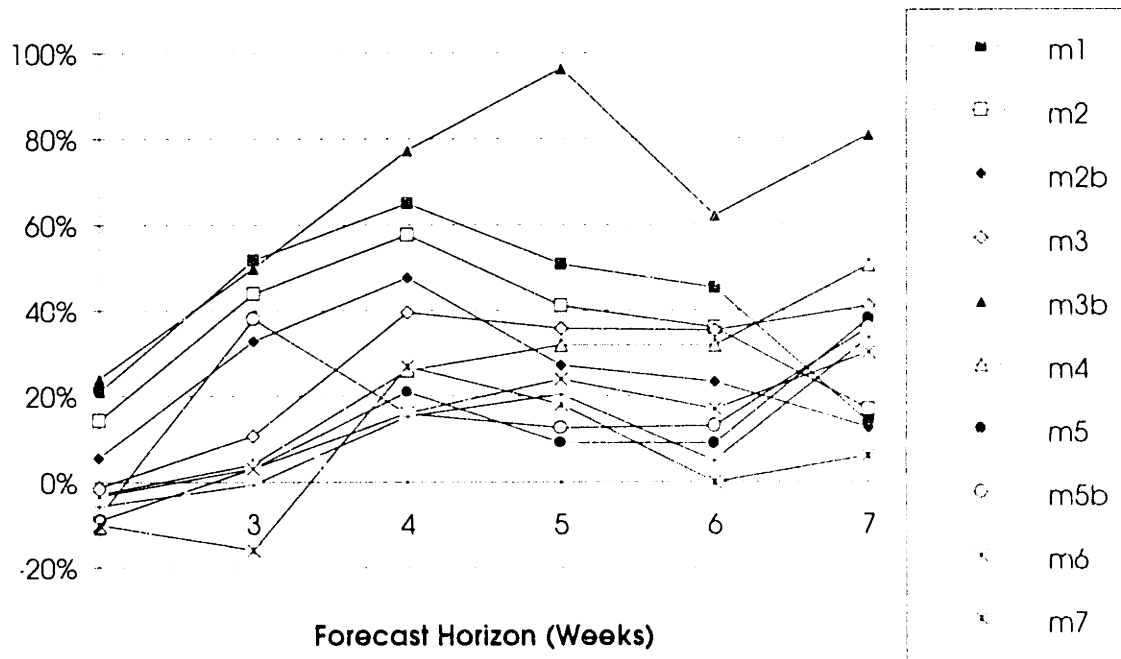


FIGURE 5.3 *MPE for Scenario 1*

As observed in Figure 5.3, most of the models display substantial positive biases over the entire range of forecast horizons. The band distribution is again

observed, although the spread between the models is much more pronounced. The magnitude of the biases of the time series models increases with the size of the forecast horizon up to 4 weeks, after which it slowly decreases as the forecast horizon further increases. The bias of regression model 3b does not conform to this behavior and becomes progressively worse as the forecast horizon increases.

The ordered performance of the models at the other end of the range of forecast horizons (7 weeks) is shown in Table 5.3

TABLE 5.3 *Theil Performance Ranking for Scenario 1: 7 Weeks Forecast Horizon*

| Model # | MPE | RMSE | MAD | MAPE | Theil |
|---------|-------|--------|--------|-------|-------|
| 7 | 0.060 | 10.284 | 8.283 | 0.517 | 0.351 |
| 6 | 0.302 | 12.125 | 8.893 | 0.323 | 0.419 |
| 5b | 0.361 | 12.292 | 9.391 | 0.715 | 0.425 |
| 4 | 0.506 | 12.921 | 10.401 | 0.793 | 0.447 |
| 7b | 0.336 | 12.947 | 9.495 | 0.676 | 0.448 |
| 5 | 0.380 | 13.573 | 9.669 | 0.793 | 0.469 |
| 3 | 0.410 | 13.719 | 8.638 | 0.861 | 0.474 |
| 2 | 0.168 | 13.886 | 11.192 | 0.786 | 0.480 |
| 2b | 0.125 | 13.951 | 11.111 | 0.800 | 0.482 |
| 1 | 0.136 | 14.569 | 10.930 | 0.645 | 0.504 |
| 3b | 0.808 | 17.672 | 13.116 | 1.018 | 0.634 |

Pickup model 7 outperforms the other models with a U value differential of approximately 20% below the next best performer (model 6). On the other hand, the top performer on week 2, has now dropped to 6th place. Model 3b again gives the poorest results. Between these two extremes, models 5b, 6, 7 and 7b are virtually indistinguishable in terms of performance, particularly when the forecast horizon is small (less than 4 weeks).

5.2.2 Scenario 2

In this scenario, the forecast horizon (FH) is maintained at 4 weeks and the size of the historical data set is varied from 4 to 10 weeks for the booking class B.

Model 3b has been eliminated on the basis of its poor performance and the realization that the relationship between the particular higher class bookings and the bookings in the discount class is not significant. Figure 5.4 shows the Theil's Inequality Coefficient results for the remaining 10 models.

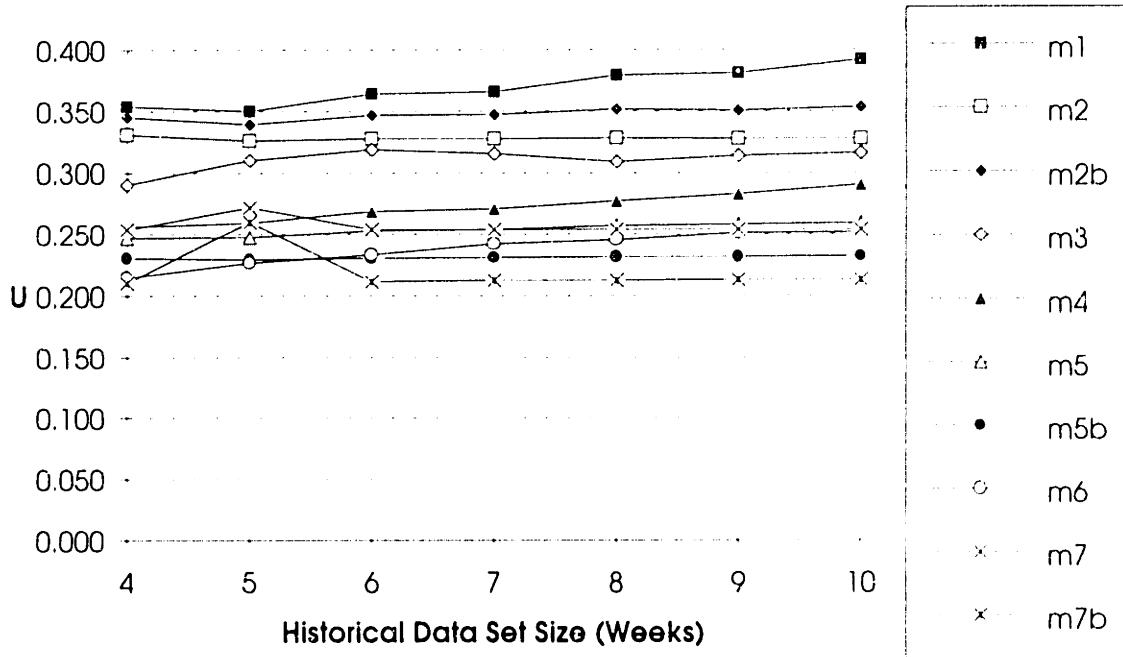


FIGURE 5.4 *Theil's Inequality Coefficient For Scenario 2*

Once again, the distinct band separation is observed. The lower band comprises the pickup models and the other models are found in the upper band. These performance bands are separated by an approximate 25% differential and the overall spread of the U values is between 0.2 and 0.4 over the entire range of the historical data set size. The U values in the lower band of pickup models are relatively constant except for models 4 and 6 whose U values shows steady growth. The upper band displays a bit more variation where the spread of the U values increases with the size of the HDS. This suggests that, in general, increasing the HDS size does not improve the performance of the models in this scenario. The

performance of models 1 and 4 deteriorates with HDS size, while model 2 displays marginal improvement over the range.

Except for the case with 5 weeks of historical data, model 7b consistently outperforms the other models. The U value differential between model 7b and the next best performer (model 5b) is approximately 8%. After 6 weeks of historical data, the U values for models 7, 7b and 5b remain perfectly constant.

All of the models are positively biased (Figure 5.5) where, apart from models 7 and 7b, the magnitude of the biases increases with the size of the HDS.

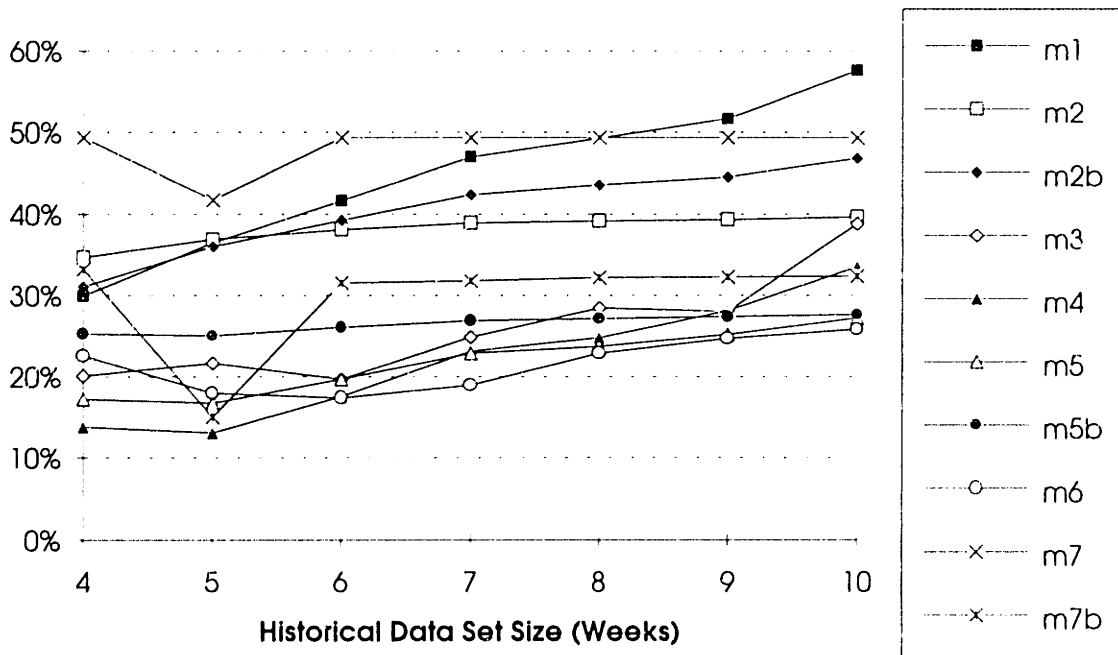


FIGURE 5.5 MPE for Scenario 2

Tables 5.4 & 5.5, show the ordered performance metrics for the extremes of the historical data set size (week 4 & week 10).

TABLE 5.4 *Theil Performance Ranking for Scenario 2: 4 Weeks of Historical Data*

| Model # | MPE | RMSE | MAD | MAPE | Theil |
|---------|-------|-------|-------|-------|-------|
| 7b | 0.332 | 5.319 | 3.720 | 0.570 | 0.210 |
| 6 | 0.226 | 5.430 | 3.767 | 0.480 | 0.215 |
| 5b | 0.252 | 5.821 | 3.933 | 0.518 | 0.230 |
| 5 | 0.172 | 6.237 | 4.138 | 0.460 | 0.247 |
| 7 | 0.494 | 6.425 | 4.515 | 0.752 | 0.254 |
| 4 | 0.138 | 6.478 | 4.274 | 0.435 | 0.256 |
| 3 | 0.201 | 7.345 | 4.766 | 0.499 | 0.290 |
| 2b | 0.347 | 8.355 | 5.734 | 0.657 | 0.331 |
| 2 | 0.311 | 8.714 | 5.900 | 0.623 | 0.345 |
| 1 | 0.299 | 8.948 | 6.078 | 0.617 | 0.354 |

TABLE 5.5 *Theil Performance Ranking for Scenario 2: 10 Weeks of Historical Data*

| Model # | MPE | RMSE | MAD | MAPE | Theil |
|---------|-------|-------|-------|-------|-------|
| 7b | 0.324 | 5.384 | 3.635 | 0.550 | 0.213 |
| 5b | 0.276 | 5.863 | 3.873 | 0.526 | 0.232 |
| 6 | 0.258 | 6.360 | 4.017 | 0.484 | 0.252 |
| 7 | 0.494 | 6.423 | 4.514 | 0.752 | 0.254 |
| 5 | 0.272 | 6.561 | 4.284 | 0.508 | 0.260 |
| 4 | 0.335 | 7.344 | 4.869 | 0.551 | 0.291 |
| 3 | 0.388 | 7.989 | 5.349 | 0.596 | 0.316 |
| 2b | 0.396 | 8.284 | 5.740 | 0.682 | 0.328 |
| 2 | 0.468 | 8.933 | 6.395 | 0.732 | 0.353 |
| 1 | 0.577 | 9.907 | 7.088 | 0.824 | 0.392 |

The rankings remain fairly constant over the entire range where the pickup models consistently outperform the regression and time series models—the advanced pickup models being the top performers with average MAD values centered around 4 (Figure 5.6).




FIGURE 5.6 *MAD for Scenario 2*

The spread of the MAD for the regression and time series models in the upper band increases gently with HDS size.

5.2.3 Scenario 3

In this scenario, the size of the historical data set (HDS) is fixed at 8 weeks and the forecast horizon (FH) is varied from 2 to 7 weeks for the booking class A.

On the basis of the performance in the discount classes the time series models were removed from model set and the focus directed towards the pickup and regression models in the remaining scenarios. This decision was based on the assumption that in the higher booking class, where the data displayed no apparent trends, the performance of the time series models would deteriorate further relative to the regression and pickup models.

Figure 5.7 shows the Theil's Inequality Coefficients for the reduced set of 7 models.




FIGURE 5.7 *Theil's Inequality Coefficient for Scenario 3*

The overall characteristic is a shallow S-curve with the U values showing little spread among the models over the range of forecast horizons. Between 2 weeks and 5 weeks out the U value increases relatively slowly with FH, after which the magnitude of U increases by a factor of 2 (1.6 for model 3) and then appears to level off. In this scenario, the regression model 3 consistently outperforms the other pickup models although for FH less than 5 weeks the difference is marginal. Among the pickup models, there is no single model that consistently performs best over the entire range.

The plot of the MPE (Figure 5.8) reveals that the biases have a strong relationship with FH.

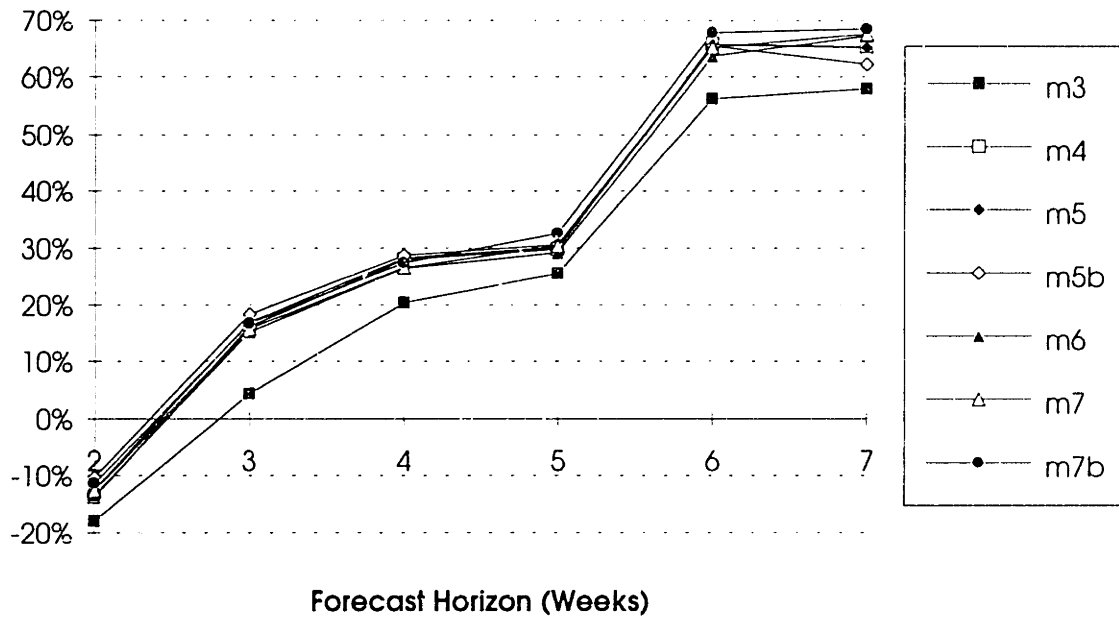


FIGURE 5.8 MPE for Scenario 3




FIGURE 5.9 MAD for Scenario 3

For forecast horizons between 2 and 6 weeks out, the magnitude of the biases behave almost linearly with FH after which they become relatively constant. Over this range, there is an approximate 50% differential between the magnitudes of the biases of the regression and pickup models.

In the case of the MAD (Figure 5.9), however, the values for regression model exceed the pickup models over the entire range of FH. This differential begins at approximately 0.5 units at 2 weeks out and diverges to approximately 1 unit at a forecast horizon of 7 weeks. The values for the pickup models display relatively little spread and remain fairly constant at 2 units between forecast horizons of 2 to 5 weeks after which a step increase of approximately 1.5 units is observed.

TABLE 5.6 *Theil Performance Ranking for Scenario 3: 2 Week Forecast Horizon*

| Model # | MPE | RMSE | MAD | MAPE | Theil |
|---------|--------|-------|-------|-------|-------|
| 6 | -0.135 | 3.490 | 1.927 | 0.349 | 0.327 |
| 7 | -0.126 | 3.546 | 1.954 | 0.536 | 0.332 |
| 7b | -0.114 | 3.654 | 2.028 | 0.540 | 0.342 |
| 5 | -0.124 | 3.655 | 1.978 | 0.555 | 0.342 |
| 4 | -0.138 | 3.659 | 1.997 | 0.658 | 0.343 |
| 5b | -0.103 | 3.718 | 1.971 | 0.572 | 0.348 |
| 3 | -0.180 | 4.445 | 2.377 | 0.432 | 0.354 |

TABLE 5.7 *Theil Performance Ranking for Scenario 3: 7 Week Forecast Horizon*

| Model # | MPE | RMSE | MAD | MAPE | Theil |
|---------|-------|-------|-------|-------|-------|
| 5b | 0.623 | 6.494 | 3.318 | 1.428 | 1.180 |
| 5 | 0.652 | 6.628 | 3.450 | 1.456 | 1.204 |
| 4 | 0.653 | 6.675 | 3.508 | 1.476 | 1.213 |
| 6 | 0.672 | 6.837 | 3.513 | 1.522 | 1.242 |
| 7 | 0.676 | 6.926 | 3.503 | 1.525 | 1.258 |
| 7b | 0.685 | 7.017 | 3.516 | 1.563 | 1.275 |

Tables 5.6 & 5.7 display the ordered performance metrics at 2 and 7 weeks out for scenario 3. The values underscore the fact that for any given metric, the spread among the pickup models is negligible resulting in the distinct streamlined characteristics relative to the discount scenarios. In addition, when compared to the discount cases, the majority of the metrics in class A are consistently higher.

5.2.4 Scenario 4

In this scenario, the forecast horizon (FH) is held at 4 weeks and the historical data set varied between 4 and 10 weeks for the booking class A.




FIGURE 5.10 *Theil's Inequality Coefficient for Scenario 4*

As seen in Figure 5.10, there is a distinct difference in the U value variation between the regression and pickup models. The U value for the regression model decreases almost linearly with HDS size while the values for the pickup models, after an initially mild divergence, remain relatively constant—centered on approximately 0.7. This suggests that increasing the HDS size has a significant

effect on the performance of the regression model while only a marginal effect on the pickup models, and virtually no effect on models 5b and 7b. In this scenario, model 6 displays the best performance with respect to U—although it is only incrementally better than the performances of models 4 and 7.

The behavior of the MAD is almost identical to the U value where the regression model shows a downward sloping characteristic while the pickup values remain relatively constant centered on 2.5 bookings.




FIGURE 5.11 MAD for Scenario 4

A closer look at the MAD values for the pickup models reveals that only model 7b remains truly constant (at 2.6) while the values for the other models decrease as the size of the HDS increases from 4 to 7 weeks after which the MAD values remain fairly constant at approximately 2 bookings—23% less than model 7b.

The MPE variation is a quite similar of the MAD behavior as seen in Figure 5.12.




FIGURE 5.12 MPE for Scenario 4

The bias of the pickup models is centered around 40% and remains fairly constant over the entire range of HDS size. The bias of the regression model decreases steadily across the entire range with over a 50% drop as the HDS size increases from 4 to 10 weeks. With respect to MPE, model 6 displays the least bias over the entire range of HDS size. As seen in Tables 5.8 & 5.9, the relative performance of the models is the same at both extremities (week 4 and week 10). Moreover, the relative performance remains constant over the entire range of HDS size. This is the first scenario where the rankings remain fixed over the entire range.

TABLE 5.8 Theil Performance Ranking for Scenario 4: 4 Weeks of Historical Data

| Model # | Theil | MPE | RMSE | MAD | MAPE |
|---------|-------|-------|-------|-------|-------|
| 6 | 0.656 | 0.365 | 4.925 | 2.448 | 0.722 |
| 4 | 0.662 | 0.414 | 4.977 | 2.354 | 0.758 |
| 7 | 0.670 | 0.371 | 5.033 | 2.486 | 0.712 |

TABLE 5.8 *Theil Performance Ranking for Scenario 4: 4 Weeks of Historical Data*

| Model # | Theil | MPE | RMSE | MAD | MAPE |
|---------|-------|-------|--------|-------|-------|
| 5 | 0.686 | 0.431 | 5.153 | 2.426 | 0.770 |
| 7b | 0.712 | 0.386 | 5.348 | 2.620 | 0.704 |
| 5b | 0.745 | 0.473 | 5.593 | 2.620 | 0.799 |
| 3 | 1.381 | 0.932 | 13.051 | 7.285 | 1.729 |

TABLE 5.9 *Theil Performance Ranking for Scenario 4: 10 Weeks of historical Data*

| Model # | Theil | MPE | RMSE | MAD | MAPE |
|---------|-------|-------|-------|-------|-------|
| 6 | 0.574 | 0.338 | 4.313 | 2.148 | 0.652 |
| 4 | 0.596 | 0.350 | 4.475 | 2.179 | 0.680 |
| 7 | 0.620 | 0.353 | 4.656 | 2.290 | 0.663 |
| 5 | 0.648 | 0.400 | 4.872 | 2.311 | 0.716 |
| 7b | 0.695 | 0.379 | 5.224 | 2.559 | 0.686 |
| 5b | 0.734 | 0.465 | 5.515 | 2.575 | 0.776 |
| 3 | 0.892 | 0.417 | 7.097 | 4.032 | 1.218 |

5.3 Unconstrained Scenarios

The unconstrained scenarios comprise scenarios 5 through 8—5 and 6 focus on the discount booking class B while 7 and 8 pertain to the higher booking class A. The constrained booking data used in scenarios 1 through 4 was unconstrained using the algorithm described in Chapter 4.

5.3.1 Scenario 5

In this scenario the historical data set (HDS) is fixed at eight weeks and the forecast horizon (FH) is varied from 2 to 7 weeks for the discount booking class B.

Based on the performance in the constrained cases, it was decided to continue the study with the top performing models only, with the regression model serving as a baseline for comparison. As a result, the classical pickup model 5 and

the advanced pickup model 7 (both with exponential smoothing coefficients $\alpha=0.2$) are removed from the model set.




FIGURE 5.13 *Theil's Inequality Coefficient for Scenario 5*

Figure 5.13 shows the Theil's Inequality Coefficients for the 5 models over the entire range of forecasting horizons. The U values vary linearly with FH between 2 and 4 weeks—at which point the regression model deviates from the other models and its U value steadily increases. Between 4 and 6 weeks the U values of the pickup models decrease gently with FH as the spread among the models diverges. The slope of the variation reverses after week 6 as the U values re-converge at week 7. Model 7b consistently outperforms the other models over the entire range where the U value is on average 10% smaller than the next best performer (Model 5b). Compared to the constrained U values (Figure 5.1), the unconstrained results appear more *streamlined* with less spread among the models—although the overall range of U values is still comparable to the constrained case. In the unconstrained case, the regression model results are more consistent

with the pickup models—particularly when the forecast horizon is relatively small (less than 4 weeks).




FIGURE 5.14 *MAD for Scenario 5*

The MAD variation is relatively small for horizons less than 5 weeks. The values increase relatively slowly with FH and the spread displays a small yet steady divergence. Beyond week 5 the magnitudes of the MAD increase rapidly to approximately 20 bookings at 7 weeks out—a factor of 4 increase. This behavior closely resembles the MAD results in scenario 3 (Figure 5.8) where the similar step function characteristic, with a break point at 5 weeks, is observed. The constrained case (scenario 1) however, does not exhibit this step behavior.

The biases of the models, determined from the MPE (Figure 15), are all positive over the range of FH (except for a horizon of 2 weeks out) and increase steadily with FH up to 5 weeks out. Beyond this turning point, the slopes reverse sign and the biases now decrease with FH—re-converging when the forecast horizon reaches 7 weeks out. In addition, there is a distinct difference in the spread before and after the turning point—the post turning point spread is a factor of

two larger. The turning point for model 5b occurs earlier at a forecast horizon of 4 weeks out. This turning phenomena is in distinct contrast to the downward characteristic of the constrained case (scenario 1).




FIGURE 5.15 MPE for Scenario 5

The ordered performance metrics at the extremities of the range of FH are shown in Tables 10 & 11. Models 7b produces the best results in all but the final horizon (7 weeks out). Apart from this top position, the rankings—and consequently the relative performance—of the models vary with FH.

TABLE 5.10 Theil Performance Ranking for Scenario 5: 2 Week Forecast Horizon

| Model # | MPE | RMSE | MAD | MAPE | Theil |
|---------|--------|-------|-------|-------|-------|
| 7b | -0.042 | 6.088 | 3.586 | 0.270 | 0.206 |
| 5b | -0.063 | 6.198 | 3.564 | 0.244 | 0.210 |
| 6 | -0.006 | 6.239 | 3.806 | 0.268 | 0.211 |
| 4 | 0.003 | 6.265 | 3.869 | 0.274 | 0.212 |
| 3 | -0.011 | 6.882 | 3.558 | 0.241 | 0.233 |

TABLE 5.11 *Theil Performance Ranking for Scenario 5: Week 7*

| Model # | MPE | RMSE | MAD | MAPE | Theil |
|---------|-------|--------|--------|-------|-------|
| 4 | 0.155 | 30.829 | 18.826 | 0.673 | 0.437 |
| 6 | 0.034 | 31.448 | 18.937 | 0.635 | 0.446 |
| 7b | 0.060 | 32.481 | 19.278 | 0.691 | 0.461 |
| 5b | 0.048 | 32.863 | 19.288 | 0.675 | 0.466 |
| 3 | 0.175 | 35.405 | 20.966 | 0.801 | 0.502 |

5.3.2 Scenario 6

In this scenario the forecast horizon is maintained at 4 weeks while the historical data set size is varied from 4 to 10 weeks for the discount booking class B.


**FIGURE 5.16** *Theil's Inequality Coefficient for Scenario 6*

Figure 5.16 shows the Theil's Inequality Coefficient for the 5 models over the range of HDS size. Models 5b and 7b both display stable behavior with the U values remaining relatively constant and quite comparable over the entire range—increasing the size of the HDS has no significant impact on the perfor-

mance of these models. Model 7b, with an average U value of 0.175, outperforms the other models. Increasing the HDS size appears to have a negative effect on the performance of models 4 and 6 as indicated by the U values increasing with HDS size. The performance of the regression model however, improves with the size of the HDS.




FIGURE 5.17 MAD for Scenario 6

The MAD for models 5b and 7b also remains constant over the entire range of HDS size. The MAD for model 6 is initially comparable to those of models 5b and 7b but grows steadily with increasing HDS until week 7 at which point it stabilizes at around 4 units. On the other hand, the MAD for model 4 is originally relatively stable yet at the same break point (week 7) its performance begins to deteriorate with increasing HDS size. The MAD of the regression model appears to be oscillating about a mean value of approximately 4.75.

As illustrated in Figure 5.18, that magnitude of the biases of the models increases with HDS size. Model 7b has a relatively constant positive bias of approximately 45%—the largest among the models. Model 4 moves from having the least bias

when the HDS size is 4 weeks to having the second largest when the HDS size reaches 10 weeks. This factor of 3 deterioration is also exhibited by model 3. The increase is not as pronounced with model 5b or model 6.




FIGURE 5.18 MPE for Scenario 6

TABLE 5.12 Theil Performance Ranking for Scenario 6: 4 Weeks of Historical Data

| Model # | MPE | RMSE | MAD | STD | MAPE | Theil |
|---------|-------|-------|-------|-------|-------|-------|
| 6 | 0.320 | 4.483 | 3.342 | 3.019 | 0.524 | 0.171 |
| 5b | 0.275 | 4.561 | 3.295 | 3.187 | 0.493 | 0.174 |
| 7b | 0.454 | 4.569 | 3.491 | 2.979 | 0.650 | 0.174 |
| 4 | 0.173 | 5.294 | 3.759 | 3.768 | 0.418 | 0.202 |
| 3 | 0.202 | 6.726 | 4.391 | 5.148 | 0.483 | 0.255 |

TABLE 5.13 Theil Performance Ranking for Scenario 6: 10 Weeks of Historical Data

| Model # | MPE | RMSE | MAD | STD | MAPE | Theil |
|---------|-------|-------|-------|-------|-------|-------|
| 7b | 0.446 | 4.577 | 3.497 | 2.984 | 0.629 | 0.174 |

TABLE 5.13 *Theil Performance Ranking for Scenario 6: 10 Weeks of Historical Data*

| Model # | MPE | RMSE | MAD | STD | MAPE | Theil |
|---------|-------|-------|-------|-------|-------|-------|
| 5b | 0.308 | 4.842 | 3.433 | 3.450 | 0.511 | 0.184 |
| 6 | 0.352 | 6.109 | 4.108 | 4.570 | 0.512 | 0.233 |
| 3 | 0.387 | 6.897 | 4.924 | 4.880 | 0.589 | 0.263 |
| 4 | 0.407 | 7.607 | 5.094 | 5.709 | 0.565 | 0.290 |

5.3.3 Scenario 7

In scenario 7, the size of the historical data set (HDS) is fixed at 8 weeks and the forecast horizon (FH) is varied from 2 to 7 weeks for the higher booking class A.


**FIGURE 5.19** *Theil's Inequality Coefficient for Scenario 7*

Figure 5.19 shows the U value performance of the 5 model subset in this scenario. The overall characteristic is almost identical to that of the constrained case (scenario 3), exhibiting the same step function behavior—although the step increase is marginally less in the unconstrained case. The regression model also

outperforms the pickup models in this case and its U values are relatively more streamlined after the *break point* forecast horizon of 5 weeks.

The overall behavior of the MAD and MPE (Figures 5.20 & 5.21) also remain virtually unchanged compared to scenario 3—although the magnitudes have diminished slightly.

Consequently, it appear as though the unconstraining of the higher booking class does not have a significant impact on the performance of the models in this scenario. This is attributed to the fact that there is not much difference between the constrained and unconstrained data since very little of the original booking data in the higher class is constrained

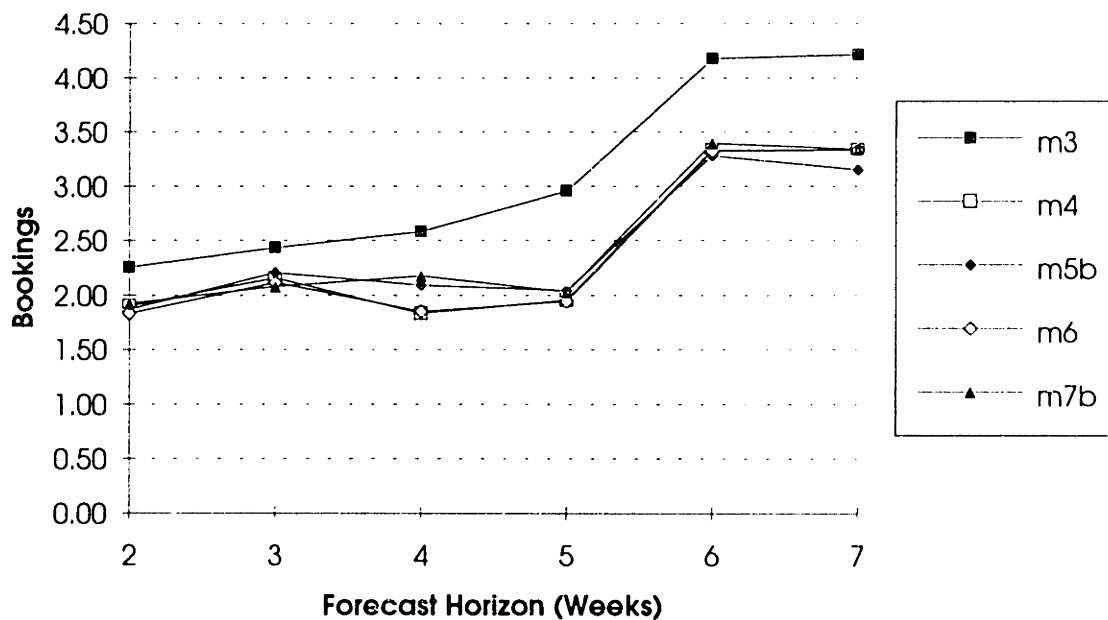


FIGURE 5.20 MAD for Scenario 7




FIGURE 5.21 MPE for Scenario 7

TABLE 5.14 Theil Performance Ranking for Scenario 7: 2 Week Forecast Horizon

| Model | Theil | MPE | RMSE | MAD | STD | MAPE |
|-------|-------|--------|-------|-------|-------|-------|
| 6 | 0.327 | -0.135 | 3.490 | 1.927 | 2.941 | 0.349 |
| 7b | 0.342 | -0.114 | 3.654 | 2.028 | 3.071 | 0.540 |
| 4 | 0.343 | -0.138 | 3.659 | 1.997 | 3.098 | 0.658 |
| 5b | 0.348 | -0.103 | 3.718 | 1.971 | 3.186 | 0.536 |
| 3 | 0.354 | -0.180 | 4.445 | 2.377 | 3.796 | 0.572 |

TABLE 5.15 Theil Performance Ranking for Scenario 7: 7 Week Forecast Horizon

| Model # | Theil | MPE | RMSE | MAD | STD | MAPE |
|---------|-------|-------|-------|-------|-------|-------|
| 3 | 1.010 | 0.551 | 6.352 | 4.230 | 5.581 | 1.395 |
| 5b | 1.180 | 0.623 | 6.494 | 3.318 | 5.641 | 1.428 |
| 4 | 1.213 | 0.653 | 6.675 | 3.508 | 5.739 | 1.476 |
| 6 | 1.242 | 0.672 | 6.837 | 3.513 | 5.928 | 1.522 |
| 7b | 1.275 | 0.685 | 7.017 | 3.516 | 6.137 | 1.563 |

5.3.4 Scenario 8

In scenario 8, the forecast horizon (FH) is fixed at 4 weeks out while the size of the historical data set (HDS) is varied from 4 to 10 weeks.

The results from this scenario are almost identical to those of the constrained case (scenario 4). Figure 5.22 shows the U value performance of the 5 models. The performance of the pickup models is almost unaffected by the size of the HDS. Models 4 and 6 display the best performance and show a gradual improvement as the size of the HDS increases while the performance models 5b and 7b remain constant. The performance of the regression model varies relatively linearly and improves significantly as the size of the HDS increases.




FIGURE 5.22 *Theil's Inequality Coefficient for Scenario 8*

When compared to the constrained case, the magnitude of the regression model bias (figure), although still considerable, has decreased. The spread of the MPE values for the pickup models has also decreased over the entire range of HDS size. The relative performance of models 6 and 4 is almost indistinguishable.

TABLE 5.16 *Theil Performance Ranking for Scenario 8: 4 Weeks of Historical Data*

| Model # | Theil | MPE | RMSE | MAD | STD | MAPE |
|---------|-------|-------|--------|-------|--------|-------|
| 6 | 0.616 | 0.365 | 4.925 | 2.448 | 4.319 | 0.722 |
| 4 | 0.623 | 0.414 | 4.977 | 2.354 | 4.431 | 0.758 |
| 7b | 0.669 | 0.386 | 5.348 | 2.620 | 4.712 | 0.704 |
| 5b | 0.700 | 0.473 | 5.593 | 2.620 | 4.994 | 0.799 |
| 3 | 1.298 | 0.932 | 13.051 | 7.285 | 10.974 | 1.729 |

TABLE 5.17 *Theil Performance Ranking for Scenario 8: 10 Weeks of Historical Data*

| Model # | Theil | MPE | RMSE | MAD | STD | MAPE |
|---------|-------|-------|-------|-------|-------|-------|
| 6 | 0.540 | 0.338 | 4.313 | 2.148 | 3.780 | 0.652 |
| 4 | 0.560 | 0.350 | 4.475 | 2.179 | 3.950 | 0.680 |
| 7b | 0.654 | 0.379 | 5.224 | 2.559 | 4.603 | 0.686 |
| 5b | 0.690 | 0.465 | 5.515 | 2.575 | 4.928 | 0.776 |
| 3 | 0.839 | 0.417 | 7.097 | 4.032 | 5.902 | 1.218 |

5.4 Observations

Based on the above results the following observations have been made:

1. Relative Performance: The pickup models outperformed the time series and regression models in the discount classes by a distinct margin (of approximately 25%) while the regression model produced the best results in the higher (full fare) booking class scenarios. Furthermore, among the pickup models, the advanced pickup models incorporating exponential smoothing with $\alpha=0.4$ consistently produced the best results.

As discussed in Chapter 1, the booking space for air transportation demand is two dimensional consisting of bookings along the *days out* axis (days before departure) as well as along the *time frame* axis (chronological order of departures). Both of these dimensions contain valuable information on the booking

characteristics of a particular flight in a particular market—the time frame dimension captures the relationship of the booking level relative to the past departures while the days out dimension provides data on the behavior of the booking history of the specific flight.

The simple time series models used in this study forecast on the basis of extrapolating the average level of the final bookings (along the time frame axis) and consequently utilize data from only one dimension of the booking space. As a result, the models do not capture the behavior of the booking profiles for individual flights. Moreover, because the variation of the final bookings is quite volatile (as seen in Chapter 4), the actual booking levels are relatively distant from the mean, resulting in substantial forecast errors. Combined, these factors account for the relatively poor performance of the time series models.

2. Forecast Horizon: In general, the performance of the models decreased as the forecasting horizon increased. This was particularly true in the case of the higher booking class where relatively small numbers of bookings were observed. When the forecast horizon is relatively small (less than 4 weeks) the performance of the advanced pickup models are indistinguishable and comparable to the results from the classical model using exponential smoothing. As the forecast horizon increased beyond 4 weeks, the advanced pickup model (employing exponential smoothing, $\alpha=0.4$) produced the best results.

As the length of the horizon increases, the location of the forecasted point moves further out in time, away from the fixed historical data set and thereby increases the chance that the booking behavior will deviate from the information contained in the historical data set. The advanced pickup models however, exploit the recent booking data from flights that have not yet departed which effectively decreases the gap between the forecast point and the historical data set, enabling local variations in the booking characteristic to be captured in the forecast—resulting in the observed superior performance. It is therefore evident that within a short-term forecast environment the local or recent data provides the most valu-

able information needed to allow a model to respond to the inherent volatility of the booking levels.

3. Historical Data Set Size: The performance of the models generally decreased with the size of the historical data set—although the performance of the pickup models incorporating exponential smoothing remained fairly stable. At first glance, this result is quite sobering. Generally in a forecasting situation, it is expected that increasing the amount of data used in the forecast model should improve performance. Nevertheless, as argued above, the nature of the short-term forecasting environment is once again brought to bear on the results. Given that the inherent variation of the actual bookings is highly volatile, it is not possible to identify overall trends. Yet if the focus is shifted to the local booking activity, *micro-trends* can be observed where the data immediately before the point of observation gives some indication of the preceding booking behavior. In light of this logic, therefore, it is the most recent data that would provide the most valuable information on the anticipated behavior of future bookings. Incorporating additional data in the historical data set introduces noise rather than adding useful information about the local booking activity. The pickup models incorporating exponential smoothing essentially extract the valuable data by weighting the recent data heaviest while suppressing the historical noise. The results for the pickup models indicate that beyond 7 weeks of historical data, the value of additional data is not significant.

When considering the performance of the regression model, the argument reverses—particularly in the higher booking class. The basis of the linear regression model uses n pairs of observations $(x_1, y_1) \dots (x_n, y_n)$ to find the least square fit to a linear relationship, where x is the bookings at a given day t and y is the final bookings. Consequently, the greater the value of n the greater the ability of the model to identify a significant relationship. This is particularly true in the case of the higher booking class where a considerable number of the paired observations

are zero. As a result, increasing the size of the historical data set improves the performance of the regression model.

4. Bias: In this study, the forecast error was defined as:

$$\text{Error} = \text{Forecast} - \text{Actual} \quad 5.1$$

The models all displayed positive biases indicating that the forecasts consistently overestimated the booking levels. The magnitude of the bias increased with the size of the historical data set. Recall that the booking data encompasses an 18 week time frame from September 1 to December 31. Over this range, the variations are mostly positive in sign and sufficiently large to shift the mean booking level above the majority of the actual bookings. As a result, because the basis for all of the models depends on the mean booking level to some degree, the forecasts are generally high, giving rise to the positive bias. As the historical data set increases, more of the variation in the demand is captured and the mean level is displaced further upward. With the advanced pickup models, however, the focus remains on the recent data and the impact of introducing additional variation is not as significant. Therefore the advanced pickup models appear less susceptible to individual flights with unusual booking activity and the bias remains stable.

Yet the magnitude of the bias is relatively quite substantial. This is attributed to the advanced pickup models being sensitive to periods with drastic changes in the booking activity. The classical pickup models react very quickly to the distortions introduced by these changes but this reaction is short-lived and the distortions are spread over fewer subsequent forecasts. The reason for this rapid reaction is because the classical pickup model counts all of the distortions at the same time. The advanced pickup models, however, spread the distortions over a greater number of the subsequent forecasts. As a result, the impact of periods with drastic changes in booking activity is felt for a relatively longer duration. This is particularly true in the case of the unconstrained scenarios where the drastic changes are more pronounced.

In the case of the unconstrained discount scenario, the magnitude of the biases increases with the forecast horizon up to a particular turning point, beyond which the biases decrease as the forecast horizon increases. This turning phenomenon is attributed to the characteristics of the unconstrained booking data. The instances of constraining occur predominantly around the holiday season which is located towards the end of the time frame, beginning around week 13. Unconstraining therefore, raises the levels of the booking profiles towards the end of the time frame (Figure 5.23).




FIGURE 5.23 Comparison of Constrained and Unconstrained Final Bookings for a Sample Market

Once the horizon moves into this region of the time frame, the actual booking levels become more consistent with the mean level and the forecast error diminishes.

This turning phenomena is not observed in the higher booking class as the effects of constraining are not as significant.

5. Effects of Unconstraining: Unconstraining did not have any significant impact on the booking characteristics of the higher booking class. This is attributed to the fact that only a small percentage of the original data in this class was

initially constrained by booking limits. As a result, the original booking data was representative of the true demand for this booking class. On the other hand, significant differences were observed when the discount bookings were unconstrained: The booking profiles—which were originally truncated—were transformed into growing exponentials. In addition, the unconstrained booking levels were generally much higher towards the end of the 18 week time frame (as shown in Figure 5.23)—attributed to the inflated demand due to the holiday period. In the unconstrained scenarios, the spread of the performance metrics among the various models decreased resulting in the similar (*streamlined*) appearance of the performance plots. In general the performance of the models improved with respect to the Theil's Inequality Coefficient—although there was a noticeable increase in the inherent bias (explained above).

5.5 Summary of Results

The following is a summary of the main findings from the results:

1. For forecast horizons less than 4 weeks, the relative performance of the pickup models is indistinguishable. The advance pickup model incorporating exponential smoothing with $\alpha=0.4$ produces the best results as the forecast horizon increases.
2. The performance of all of the models decreased with forecast horizon.
3. Increasing the size of the historical data set beyond seven weeks did not have a significant impact on the performance of the models.
4. The models all displayed an inherent positive bias.
5. Unconstraining the booking data improved the performance of the models with respect to the Theil's Inequality Coefficient, yet the magnitude of the inherent bias increased.

Chapter 6

Conclusions

6.1 Research Findings

The findings drawn from this study pertain to (1) the nature of the short-term forecasting environment and (2) the performance of the selected models in this environment.

6.1.1 The Short-Term Forecasting Environment

In this study, the short-term forecasting environment was defined by confining the forecasting horizon to 8 weeks out and restricting the size of the historical data set to 10 weekly departures. The resulting time frame became sufficiently small such that, although there was considerable variation in the weekly booking levels, no underlying trends were apparent—the variation appeared purely stochastic and highly volatile. Furthermore, because seasonal and cyclical trends were not easily identified in this environment, forecasting models which relied on the extrapolation of trends produced relatively poor results.

Notwithstanding, one distinct seasonal variation was observed where the booking levels increased significantly in the vicinity of the holiday period. It is under these circumstances, where the levels of demand became inflated that the effects of capacity constraints were most pronounced. Consequently, it was necessary to unconstrain the bookings in order to estimate the true demand. In this study, there was a distinct difference in the characteristics of the constrained and

unconstrained discount class bookings. Unconstraining the booking data not only removed the truncated appearance of the booking profiles but also improved the performance of the forecasting models overall.

6.1.2 The Performance of the Selected Models

Based on this study, it is clear that the pickup models consistently outperformed the regression and time series models in the various scenarios. When the forecast horizon was relatively small (less than 4 weeks out), however, the performance of the pickup models was literally indistinguishable. As the forecast horizon increased, the advanced pickup model incorporating exponential smoothing produced the best results. The superior performance of the advanced pickup model was attributed to the use of the most recent data where the focus was on the *local* booking activity. Nevertheless, the advanced pickup model was found to be more sensitive to periods where the booking activity changed drastically. Although all the models are subject to the distortions created by rapid changes in demand, the advanced pickup model spread the effects of this distortion over a greater number of subsequent forecasts which resulted in a consistently larger bias.

All of the models displayed positive biases indicating that the majority of the forecasted final bookings were over estimated. This overestimation was attributed to the high mean booking levels with respect to the actual bookings due to the relatively large positive variations in demand—inherent in the short-term environment.

The observed biases could be mitigated through the introduction of a compensatory error term in the models. This becomes more of a challenge when considering the pickup models, however, as these models do not have coefficients to calibrate and compensation would therefore have to be done on a case specific basis. Because the observed bias is an average value, it would not be appropriate to simply adjust the individual forecasts by the magnitude of the bias.

Of particular interest is the discovery that, within the short-term forecast environment, increasing the size of the historical data set did not have a significant impact on the performance of the models. In the case of the pickup models, using more than 7 historical weekly departures had an adverse impact on performance. The performance of the advanced pickup models, however, remained stable throughout— independent of the size of the historical data set. This result is counter-intuitive to conventional statistical principles that suggest that the accuracy of an estimation increases with sample size. Yet, a distinction should be made between the size of the number of observations used to calibrate coefficients in the forecast models, to which this statistical theory applies, and the size of the historical data set used to generate the forecast. The pickup models do not employ coefficients or constants in the forecasting process and therefore the statistical argument is not applicable.

In light of the above conclusions, it is clear that the advanced pickup models should be considered as one of the preferred techniques for forecasting the short-term demand for air transportation. The strength of these pickup models lies in weighting the most recent data heaviest— particularly the bookings from flights that have not yet departed— coupled with the use of data from the two dimensions of the booking space: the final bookings across the time frame as well as data from the booking profiles at given days out before departure. The decomposition of the pickup interval into smaller increments also facilitates focusing on the local booking activity and has consequent improvements on performance.

Indeed, this model has certain shortcomings, yet, the computational efficiency, ease of implementation, relatively low data requirements, and proven performance warrants the further development and utilization of this model.

6.2 Revenue Impact of Forecast Errors

Studies indicate that the revenue impact of the forecast bias is strongly influenced by the overall demand level. In one particular analysis, Curry [1] studied

the impact on revenue using Monte Carlo simulations. For each set of conditions, the percent revenue achieved was computed defined as the revenue achieved with the forecast error divided by the revenue that could have been achieved with full knowledge of demand. The typical results are shown in Figure 6.1.




FIGURE 6.1 Revenue Impact vs. Forecast Error [1]

If the demand is low then the forecast errors have little impact because there are few inventory restrictions regardless of the forecast. On the other hand, the same error can have a larger impact if the demand is high because too many or too few seats would be consistently reserved. Over-forecasting will save too many seats and the flight is more likely to depart with some seats empty, thus losing the entire amount of a fare. Under-forecasting will lead to full flights but too many discount passengers onboard (not enough seats saved for the late-booking, high revenue passenger). The loss in this situation is the difference between the full

fare and the discount fare. As a result, over-forecasting has a more significant impact on revenue than under-forecasting.

6.3 Avenues for Further Study

The following represent areas for further study:

1. Given the revenue impact of over-forecasting and the positive bias inherent in the advanced pickup model, there is a need to study the effectiveness of incorporating an error factor to compensate for the bias in the pickup models. As alluded to above, this would have to be done on a case specific basis as the pickup models do not lend themselves to traditional calibration techniques.
 2. Although the advanced pickup models appear less subject to individual flights with odd booking patterns, they are susceptible to periods of odd booking activity. One possibility to mitigate this sensitivity would be to incorporate *adaptive filtering*, in the estimation of the average incremental pickup. Adaptive filtering is a variation of exponential smoothing where the magnitude of the smoothing constant α depends on the average error and the absolute error of the previous forecast.
 3. The markets in this study comprised a collection of short, medium and long hauls. Studying the performance of the models on the basis of market type would add some insight into whether the booking characteristics and consequent performance of the models are independent across market type. This would allow the potential of utilizing different techniques for different market types to be addressed.
-

Bibliography

- [1] Curry, R.E. *Forecasting for Revenue Management*: Technical Brief, The Revenue Management Quarterly Scorecard, 1994.
- [2] Belobaba, Peter *Air Travel Demand and Airline Seat Inventory Management*. MIT Thesis: MIT Flight Transportation Lab, 1987.
- [3] Lee, Anthony *Airline Reservations Forecasting*: Proceedings of the 28th Annual AGIFORS Symposium, 1988.
- [4] Barry C. Smith, John F. Leimkuhler and Ross M. Darrow, *Yield Management at American Airlines*: Interfaces, 1992.
- [5] Lee, Anthony *Airline Reservations forecasting*: MIT Thesis: Flight Transportation Lab, 1988.
- [6] Simpson, R.W. *Notes for Air Transportation*. MIT Flight Transportation Laboratory Lecture Notes: Massachusetts Institute of Technology, 1982.
- [7] Doganis, *Flying Off Course*: The Economics of International Airlines. Harper-Collins, 1991.
- [8] Belobaba, Peter *The Revenue Enhancement Potential of Airline Yield Management Systems*: Proceedings of the Advanced Software Technology for Air Transport International Conference & Exhibition, 1992.
-

- [9] Taneja, N. K. *Airline Traffic Forecasting*: Lexington Books, 1978.
- [10] Benton, W. K. *Forecasting for Management*. Addison-Wesley, 1972.
- [11] Montgomery, Douglas C., and Johnson, Lynwood A. *Forecasting and Time Series Analysis*: McGraw-Hill 1976.
- [12] Box, G., and Jenkins, G., *Time Series Analysis—Forecasting and Control*: Holden-Day
- [13] Goodell, Brown, Robert *Smoothing Forecasting and Prediction of discrete Time Series*: Prentice Hall 1963.
- [14] Jenkins, G. M. , and Watts, D. G. , *Spectral Analysis and Its Applications*. Holden-Day, 1968.
- [15] Anderson, T. W., *The Statistical Analysis of Time Series*. Wiley, 1971.
- [16] Granger, C. W. J., *Spectral Analysis of Economic Time Series*. Princeton University Press, 1964.
- [17] Henning, D., *Air Traffic Forecasting Methods*: Airport Forum 6, No. 2, 1976.
- [18] Garvett, D. S., and Tenaja, N. K., *New Directions for Forecasting Air Travel Passenger Demand*, MIT Flight Transportation Laboratory Report No. R74-3, 1974.
- [19] Kanafani, A.K. *Transportation Demand Analysis*: McGraw-Hill, 1983.
- [20] Ben-Akiva, Moshe, and Lerman, *Discrete Choice Analysis: Theory and Application to Travel Demand*: MIT Press, 1985.
- [21] Littlewood, Kenneth *Forecasting and Control of Passenger Bookings*.: Proceedings of the 12th Annual AGIFORS Symposium, 1972.
- [22] Scandinavian Airlines System. *SHOPS: A Short-Term Passenger Forecasting System*, 1978.
-

-
- [23] Duncanson, Archie V. *Short-Term Traffic Forecasting*: Proceedings of the 14th Annual AGIFORS Symposium 1974.
- [24] Sa, Joao *Reservations Forecasting In Airline Yield Mangement*. MIT Thesis: Flight Transportation Lab, 1987.
- [25] Brummer, Mark et al. *Determination of Potential Load Capacity Based on Observed Load Factor Data*: A study for Northwest Airlines: St. Olaf College Undergraduate Practicum Group, 1988.
- [26] Ben-Akiva, Moshe, et. al. *Improving Airline Passenger Forecasts Using Reservation Data.*: Presentation at Fall ORSA/TIMS Conference, St. Louis, MO, 1987.
- [27] Rothstein, M. *An Airline Overbooking Model*: Transportation Science, 1971.
- [28] Maddala, G.S. *Limited-Dependent and Qualitative Variables in Econometrics.*: Cambridge University Press, 1983.
- [29] Belobaba, Peter *Yield Management Forecasting Made Simple*: Presentation to 4th International IATA Yield Management Conference, 1992.
- [31] Adams, Warren and Vodicka, Michael *Short-Term Forecasting of Passenger Demand and Some Applications in Qantas*: Proceedings of the 27th Annual AGIFORS Symposium 1987.
- [32] L'Heureux, Ed *A New Twist In Forecasting Short-Term Passenger Pickup*: Proceedings of the 26th Annual AGIFORS Symposium 1986.
- [33] Hogg, Robert, and Ledolter, Johannes *Applied Statistics for Engineers and Physical Scientists*: Macmillan 1992.
- [34] Belsey, D.A. et. al. *Regression Diagnostics: Identifying Influential Data and Sources of Collinearity.* : Wiley, 1980.
- [35] Larsen, R.J. and Marx, M.L. *An Introduction to Mathematical Statistics and its Applications*: Prentice-Hall, 1986.
-

[36] Amemiya, T. *Regression Analysis When the Dependent Variable is Truncated Normal*. *Econometrica*, Volume 41, p. 997, 1973.

[37] Belobaba, Peter *TWA Reservations Analysis: Demand Distribution Patterns*. MIT Flight Transportation Lab Report. : Massachusetts Institute of Technology, 1985.
