FLIGHT TRANSPORTATION LABORATORY
REPORT R98-4

FORECASTING FOR AIRLINE NETWORK REVENUE MANAGEMENT; REVENUE AND COMPETITIVE IMPACTS

BY: JEFFREY S. ZICKUS

DEPARTMENT OF
AERONAUTICS AND ASTRONAUTICS
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
CAMBRIDGE, MASSACHUSETTS 02139-4307

FLIGHT TRANSPORTATION LABORATORY
REPORT R98-4

FORECASTING FOR AIRLINE NETWORK REVENUE MANAGEMENT; REVENUE AND COMPETITIVE IMPACTS

BY: JEFFREY S. ZICKUS

Forecasting for Airline Network Revenue Management: Revenue and Competitive Impacts

by

Jeffrey S. Zickus

Submitted to the Department of Civil and Environmental Engineering
on May 15, 1998, in Partial Fulfillment of the Requirements of the Degree of
Master of Science in Transportation

Abstract

Airline revenue management entails protecting enough seats for late-booking, high-fare passengers while still selling seats which would have otherwise gone empty at discounted fares to earlier-booking customers. In the evolution of revenue management to network origin-destination control, previous research has shown that revenue gains of some seat optimization algorithms can be much lower than others. One possible reason is the process by which demand estimates are generated; namely, forecasting and detruncation. Forecasting is used to estimate passenger demand based on historical flight data, while detruncation makes projections of what demand would have been in cases where the historical data has been constrained by a capacity limitation. This thesis explores the question of the interaction between forecasting methods, detruncation methods, and seat optimization algorithms on a simulated airline network, using the Passenger Origin-Destination Simulator (PODS) revenue management simulation tool, which models a network environment with two competing airlines.

Changes in the forecasting and detruncation methods in combination with the seat optimization algorithms were tested in order to see what revenue impacts resulted. Additionally, passenger loads, forecasts, and fare class availability were examined to understand the reasons behind the observed revenue results. The simulations showed that seat optimizers which had relatively poor performance using a standard forecasting and detruncation method had substantial revenue increases when different forecasting and detruncation combinations were implemented. The results also indicate that the better combination of forecasting and detruncation causes higher revenues for all seat optimization methods tested, as a better passenger mix is realized due to higher levels of detruncation and more accurate forecasts. However, the sensitivity of the seat optimizers to the forecasting and detruncation methods remains mixed. Inferior detruncation (or forecasting) methods on a network can offset the revenue gains resulting from improvement to origin-destination control from leg-based control for some seat optimization algorithms.

Thesis Supervisor: Dr. Peter P. Belobaba
Title: Associate Professor of Aeronautics and Astronautics

Table of Contents

List of Figures

List of Tables

Room 14-0551
77 Massachusetts Avenue
Cambridge, MA 02139
Ph: 617.253.5668 Fax: 617.253.1690
Email: docs@mit.edu
http://libraries.mit.edu/docs

Page 12 is missing from original

Chapter 1

Introduction

1.1 Forecasting and detruncation applied to airline revenue management

Revenue management (more commonly referred to as yield management) has become a subject of heavy research and has undergone widespread implementation over the past fifteen years by many of the world's airlines. In an airline context, it symbolizes the desire by an air carrier to obtain the highest amount of revenue possible from the traveling consumer by enticing them to pay a fare commensurate with their economic willingness to pay (WTP). In earlier years of government-regulated air travel in the United States, this was not possible, as ticket prices between origins and destinations were fixed by the Civil Aeronautics Board. However, with the advent of airline deregulation in 1978, pioneering firms in the industry have been clever about how they go about performing such a task--various restrictions are placed on different fare products in order to allow for price discrimination and better market segmentation. These restrictions usually include (but are not limited to) provisions such as advance purchase requirements, length of stay requirements, and refundability restrictions. For example, a leisure passenger flying from Boston to Chicago may be able to find a non-refundable, 21-day advance purchase round-trip ticket with a Saturday night minimum stay for as low as \$150, while a business passenger purchasing a ticket one day in advance with no minimum stay restrictions and full refundability may pay as much as \$1000 for travel in the same market.

From the airline's business perspective, they would ideally like to fill the seats on the aircraft with as many of the high-fare customers as possible (i.e., those with the highest WTPs), and then offer the remaining empty seats at more discounted fares to attract customers with lower willingness-to-pay. The primary complication in air travel stems from the fact that the lower one's WTP, the earlier the reservation tends to be made. So, while airlines will find it less profitable to accept all low-fare customers initially, turning them all away in hopes of filling the plane with late-booking, full-fare business customers will also tend to reduce loads and decrease profits (with the exception of cases with extremely high-demand). The essence of revenue management is therefore twofold. First, *differential pricing*, which entails the determination of the set of fare products to be offered and their associated restrictions; and second, *seat inventory control*, which entails how many seats to make available for each of these fare products[1]. Given that we have a set of fare products and an established pricing regime, revenue management comes down to the determination of seat protection limits and capacities to a set of fare classes in order to maximize the potential revenue of the airline by reserving and selling as many high-fare seats as possible while filling otherwise empty seats by offering lower-fare tickets with travel restrictions. (As long as the fare or marginal revenue of a low-fare passenger exceeds the marginal cost of the available empty seat on the aircraft, then the airline should theoretically be willing to take the passenger). A good and complete discussion of the principles of revenue management applied to the airline industry can be found in Belobaba[2], while several other sources also provide detailed discussions; namely, Bodily and Weatherford[3], Smith[4], Williamson[5], and Wilson[6].

The goal of any airline practicing some form of revenue management is to attempt to segment the market demand by economic willingness-to-pay, and then offer a different fare product to each of these customer segments. To maximize revenue on a single flight leg, an "ideal" combination of fare class booking limits is determined, based upon projected estimates of the potential demand by fare class on that flight leg or path and the given pricing structure. Doing this is relatively straightforward; however, a major obstacle arises when the scope is expanded to consider the massive hub-and-spoke networks of any large airline where thousands of different O-D market

[1] Belobaba (1992). The first section of the article gives a brief introduction into how this is done.
[2] Belobaba (1987). Chapter 1 analyzes the prices and products offered by airlines in the competitive environment, Chapter 2 looks at demand from the consumer perspective, and Chapter 3 provides a thorough discussion of the seat inventory control background and process.
[3] Bodily and Weatherford (1992). A general description of perishable-asset revenue management is given along with a classification of the elements involved in such problems; previous research on some of these elements is also reviewed.
[4] Smith et al. (1992). Component descriptions of yield management and how it was applied at American Airlines are analyzed.
[5] Williamson (1992). Section 2.1 describes the basics of seat inventory control, while Section 2.2 illustrates its application to an airline network.
[6] Wilson (1995). Chapter 1 describes good motivation for practicing revenue management, while Chapter 2 details different approaches to the problem.

combinations are connected by hundreds of different flight legs and several different fare classes. In these cases it could be more profitable to carry lower-fare class connecting passengers than higher-fare class local ones, if the connection passenger's absolute fare value is higher; although if two local passengers could be accommodated instead of a single connecting passenger, then the two local passengers will generally contribute more revenue than the single connecting one. This is the O-D seat inventory control problem, explained in more detail by Belobaba[7]. Solutions to this problem would be evident through network optimization, although there are several obstacles to achieving this goal, primarily that of the sheer size of the problem, given the large number of flights and cities served by most major airlines. In fact, Belobaba also gives an example where more than 2.5 million seat inventory levels can exist for even a medium-sized US airline[8]. Difficulty therefore arises in trying to efficiently optimize revenues over the entire *network*, (i.e., how many seats should be protected for each fare class on each flight leg or path in the network?).

As stated earlier, the higher fare-paying customers (primarily business travelers) are usually those who make reservations much later in the booking process than the lower fare-paying ones (generally leisure travelers). Therefore, when practicing revenue management, the airline must make an initial guess as to how many seats to offer to the early-booking, low-fare customers and how many to reserve for the high-yield latecomers. These booking limits are in fact dynamic, as they can then be altered as time goes along and bookings come in (i.e., the airline begins to have *actual* information about the flight as opposed to just *forecasted* information), but good initial estimates are needed to keep from selling out the aircraft early in the process with too many low-fare passengers and then later turning away high-fare demand. So the revenue management question ultimately translates to, "How many seats should be made available to each fare class on each flight at any given point in time prior to departure, in order to maximize expected network revenues?" It should be noted that such a question is addressed assuming a static schedule, fixed

---

[7] Belobaba (1995). Section 1.2 gives more detailed examples of this problem, illustrating how network revenue maximization is not necessarily obtained from maximization of each of the flight leg revenues individually.
[8] Belobaba (1992). An example illustrating the need for computerized algorithms to solve these problems is given in the second section of the report.

aircraft capacities, and a given pricing regime, along with an estimate of what the potential customer demand for travel is in any given O-D market.

In the ideal case, we would like a complete optimization of not only fare class limits, but simultaneously of schedules, fares, and O-D markets served. However, such a feat, if possible (it would be extremely difficult to predict competitors' responses to changes by a given airline--so any such optimization done by an airline would therefore have to make assumptions about competitive responses), would entail an optimization scheme of billions of equations and decision variables. So, given that we are constrained to the optimization of fare class booking limits, how do we determine the market demand upon which our protection limits are based?

This is precisely where methods for forecasting and detruncation play an important role in the yield management process, as an airline does not have a clear idea on the present day of what the exact demand for a particular flight will be on some future day. But, to effectively use their revenue management tools, relatively accurate predictions of demand are necessary. Attempts must therefore be made to make some conjecture of this future demand, using any information we already have on similar flights which have already departed and/or demand as of the current day for similar flights on future days. A forecasting method does just that--it provides us with a projected demand by fare class for a given flight, based on complete historical observations of similar flights and incomplete current observations for future flights. Different forecasting methodologies can be used to obtain such an estimate by transforming this data in different ways, using some or all of the aforementioned information.

Given that we can use our forecaster to obtain an estimate of demands on different flights, and given our pricing regime, operating schedules, and aircraft capacities, we should now be able to obtain a beneficial allocation of seats to the different fare classes for each flight in the system as to maximize network revenues. However, one other slight complication arises when we consider these demands for airline flights. Airplane capacities are more or less fixed on a given flight by aircraft type and schedule rotation constraints. Hence, once a fare class on a future flight becomes full (i.e., the number of bookings reaches the total number of seats on the aircraft, assuming no overbooking), the data obtained no longer paints an accurate picture of what the actual demand would have been for the flight if no capacity restriction had been in place. Therefore, detruncation of the data adjusts for this by projecting an estimate of what *unconstrained* demand would have been based on historical bookings (i.e., what would the demand have been given no capacity restriction?). Customers making reservations and desiring a particular ODF (origin, destination, fare class) are only given the best *available* options; if their preferences are not met, they are subsequently left to decide if they wish to "sell-up" to a higher fare class or if they wish to reject all offerings and travel on another carrier. In current airline practice, this information is not recorded in the airline database, so some projection of it must be inferred from the data using a statistical detruncation method.

1.2 Motivation

There are two primary reasons for the experiments performed in this thesis. First, prior studies have shown that the application of different seat optimization algorithms prior to and during the reservations process can provide revenue gains on the order of 8-10% for an airline in a competitive market where the competition does not use any form of revenue management, and even as much as 1-2% in a simulated network situation where both carriers were using some form of revenue management (see Wilson[9] and Lee[10] for a discussion of different yield management case results on a single leg basis and a network basis, respectively). But most of the prior simulations have used identical forecasting and detruncation methods for each yield management algorithm; few, if any, have tested the revenue impacts of changes in forecasting or detruncation methods in conjunction with different seat optimization cases. Therefore, it is of interest to study the impacts of changes (by one or both airlines) in these forecasting and detruncation methods in competitive scenarios under different combinations of yield management methods. This will allow us to make conclusions about whether the gains from implementation of a better yield management system can fully be attributed to the seat optimizer itself, or if a forecasting or detruncation method was possibly hindering or helping the seat optimization algorithm and biasing the yield management performance results.

---

[9] Wilson (1995). Section 5.1 tests the model of two identical competing carriers (only the seat optimizer can be changed) serving one isolated O-D market, each with one daily departure.

[10] Lee (1998). Section 4.1 provides motivation, setup, and results for network-based O-D control methods.

Second, of previous research which has tested different forecasting and detruncation methods, little has been applied in combination with seat optimization routines to a large network scenario. Skwarek[11] analyzed and examined changes in the forecasting and detruncation methods on a single flight leg; however, more effort was concentrated on analyzing the relative performance of the methods themselves in combination with a single seat optimization routine. The simulation runs performed were for two competing airlines in a single O-D market case, where one airline had the flexibility of changing the forecasting or detruncation algorithm, but no tests were performed to analyze the impacts of variation in the seat optimization algorithm. In the domestic US airline industry; however, a much more complex network structure exists, where several hub-and-spoke networks are joined by numerous connecting hub airports (each generally operated by an individual air carrier). This gives passengers the ability to fly from various origins to various destinations, usually by making a connection through one of these hub cities. And despite the fact that most major US airlines use some sort of seat optimization routine in their practice of yield management, the methods are by no means uniform across carriers. Also of interest, therefore, is a simulation of such network conditions, where a variety of changes are possible for each input parameter, to better simulate the real-world industry. From this, conclusions can also be made as to which forecasting/detruncation methods perform best in combination with the different seat optimization routines tested.

Why study the impacts of different forecasting methods, slight as they may be? There are several inherent problems with inaccurate demand forecasts, which will inevitably lead to less-than-optimal seat allocations and network revenues, for two primary reasons. First, an underprediction of high-yield passengers results in the underprotection of high-fare class seats and therefore revenues become diluted by excessive numbers of low-fare customers who book reservations early and end up displacing their higher-fare counterparts. Second, an overprediction of demand can cause protection of too many seats for high-fare passengers, resulting in low-fare passengers being turned away early in the booking process, while seats reserved for the late-booking, higher-fare passengers end up going empty as the expected demand does not materialize. Either way,

[11] Skwarek (1997).

some portion of potential revenues are not realized, thereby highlighting the need for methods producing as precise of a forecast as possible.

1.3 Objective of thesis

Based upon the above examples, the primary objective of this thesis is therefore to study the impacts and revenue effects of different forecasting method choices and different detruncation algorithms as applied to a simulated airline network in a revenue management context. That is, the Passenger Origin-Destination Simulator (PODS), developed at Boeing[12], will be used to simulate a variety of these different combinations in conjunction with different seat optimization algorithms. The primary purpose of this thesis is not to explore the details and relative advantages and disadvantages of a particular forecasting methodology (see Wickham[13] and Skwarek[14] for discussions of forecasting methodology and their applications to a simulated airline network) nor to compare a particular methodology's predictions with the actual realized demands, as this is very difficult to measure in the airline world.

The reasoning behind this difficulty stems from the interdependence of fare class demands. Although an airline will have an exact count at the time of departure of the number of passengers on a flight leg (and they can even obtain the individual numbers from different flight paths if desired), the capacity restrictions in place might have caused potential demand to be spilled to other airlines or resulted in sell-up to a higher fare class by a passenger. Hence, the "feedback" effect makes measurement of actual demand difficult--the lower our forecasts are for a given ODF, the fewer seats are allocated, and the more likely passengers are to be spilled; however, if passengers are likely to be spilled, then the forecast was probably too low. The problem in obtaining actual demands should be apparent. Therefore, more emphasis will be placed on the

---

[12] Hopperstad (1997). A very brief description of PODS version 6 (primarily the same as version 7, which will be used in this thesis) is provided; system architecture is described, along with forecasting and seat optimization combinations.

[13] Wickham (1995). The application of different forecasting methods applied to airline bookings was tested; Chapter 3 details the models tested, Chapter 4 explains the background and procedure, while Chapter 5 presents and analyzes the results.

[14] Skwarek (1997). Section 3.1 provides a review of former approaches to airline forecasting while Section 5.2 details the different methods to be tested and Section 6.1 provides results from those forecasting tests.

analysis of the merits and weaknesses of and interaction between commonly used forecasting, detruncation, and seat optimization methods available in yield management algorithms.

This thesis basically attempts to expound on previous research involving the PODS simulation; allowing changes in both forecasting and detruncation methods while testing different seat optimizers under a more real-world network scenario, with six individual spoke cities, two connecting airport hubs, and two competing airlines. There are then two ultimate goals. First, a robust set of simulated airline network results will be generated which endeavor to pinpoint the reasons why different seat optimization systems perform differently under different forecasting and/or detruncation methods. Second, "best-case" combinations of algorithms will be obtained to determine the most advantageous network revenue management practices for the airline under a given set of circumstances. Both of these will be done based upon analysis of different performance measures, such as system revenues, passenger loads by path/leg, fare class closure rates, and forecasted remaining demands by time frame. The first of these provides a good overall measure of relative performance among the yield management methods, while the other three give more insight into why the revenue trends are occurring.

1.4 Structure of thesis

The first chapter of the thesis provides a brief introduction to the revenue management problem for an airline, along with motivation for testing and using different combinations of forecasting and/or detruncation methods available in the simulation. Small illustrations of how forecasting is used in airline revenue management and why it is important is also presented, along with the basic objectives of the thesis. Finally, yield management-related terminology to be used throughout the thesis is defined.

The second chapter gives insight into the inner workings of the different forecasting and detruncation methods to be tested, with complete descriptions of the different methods tested in the PODS simulations and how they are used by PODS. Brief descriptions of the seat optimization routines are provided, with descriptions of the interaction these methods have with the different yield management methods presented.

The third chapter then provides a brief description of the PODS simulator and its characteristics, along with descriptions of the airline network used in the simulation runs and details of what is to be expected for variations in the PODS input parameters such as demand factor changes and competitor airline parameter changes. Differences arising from the network case used here and the single-leg case used in previous research are discussed, and several output performance metrics are described.

The fourth chapter provides a complete analysis of different combinations of forecaster, detruncation routine, and seat optimization algorithm. Different base cases are used, and results subsequently compared in a variety of manners to determine the relative advantages and disadvantages of each particular method or combination of methods. Section 4.1 analyzes the resultant impacts for Airline A against base case yield management done by Airline B under different combinations of forecasting and detruncation methods. Section 4.2 looks at the impacts of "competitive scenarios," where both airlines either match the seat optimization method or have full flexibility in changing their forecasting, detruncation, and seat optimization algorithm. Furthermore, a synopsis of which forecasting/detruncation method perform best with a given seat optimization algorithm is given. Section 4.3 analyzes the relative revenue results from changes in one of the parameters in a particular detruncation routine to determine the sensitivity of the parameter. Finally, Section 4.4 explains the factors contributing to the resultant revenue increases or decreases, and possible reasons for why various results were obtained will also be analyzed. Additionally, results will be compared with any previous ones (i.e., from earlier PODS versions), and reasons behind any similarities or differences will be explained.

The fifth and final chapter provides a brief review of the thesis methodology along with a general summary of results. Highlights of the thesis objectives are analyzed--that is, whether forecasting and detruncation act as a catalyst within some of the seat optimizers, and what the best overall combinations of forecaster, detruncator, and seat optimizer are. Finally, questions still unanswered by this research are posed, from which possible or interesting directions for new research are given.

1.5 General nomenclature

Various terminology specific to airlines and yield management will be used throughout this thesis; provided in Table 1.1 below are definitions of some of the recurring terms and notation which will be used in subsequent chapters.

Term	Definition
PODS	Passenger Origin-Destination Simulator, the simulation tool used in the experiments
Revenue management	The allocation of seats on one or more flight paths to passengers based on fare class limitations, expected demands, etc.
Path	Flight or flights which will provide travel from origin to destination
O-D	Origin and destination path
Fare class	Classification signifying fare category paid by passenger
ODF	Designation of origin/destination/fare class
Bookings	Reservations accepted by the airline
Time Frame (Booking Interval)	Time period between seat optimizations during which passenger bookings are accepted
Overbooking	Permission of bookings in excess of aircraft capacity to help account for no-shows
Fare class closure	Occurrence when the number of bookings in a fare class reaches its booking limit and no further bookings are accepted
Seat optimizer	Algorithm for determining booking limits by fare class
Forecaster	Method by which demand projections are made for path/fare classes which do not yet have complete booking information
Detruncator	Method by which flights whose path/fare class bookings reached capacity are adjusted to gain a prediction of demand had there been no capacity restriction
Demand factor	Ratio of the *average* realized demand to the aircraft capacity
Booking curve	General profile by which passengers of a specific type book their reservations
System revenue	Total revenue for all flight legs flown by the airline
Passenger pickup	Number of incremental passenger bookings for a fare class in a specific time frame

Table 1.1: Revenue management/PODS nomenclature

These will provide a framework for much of the discussion and analysis in the ensuing chapters of this thesis, as different forecasting and detruncation methods are tested and analyzed in combination with several seat optimization routines. Chapter 2 provides a background of the various methods which will be tested.

Chapter 2

Forecasting/Detruncation Methods Explained

Before the results of the simulations carried out using PODS are presented and discussed, it is helpful to gain a better understanding of how the internal algorithms operate within the PODS system. This chapter will therefore explain the inner workings of the different forecasting and detruncation methods which will be tested in the subsequent chapters. First, two different forecasting methods tested in the PODS simulations are presented--the pickup forecaster and the regression forecaster. Second, two different detruncation algorithms available in the simulation are discussed--booking curve detruncation and projection detruncation. Finally, as this thesis is a study of the passenger effects and revenue impacts of changes in these forecasting/detruncation approaches under different yield management algorithms, a brief introduction to the different types of yield management will be given, and attempts will be made at explaining the subsequent interaction which occurs between the seat optimizer and the forecasting and detruncation routines it uses.

2.1 Forecasting methods

A forecast can be defined as "a quantitative estimate (or set of estimates) about the *likelihood* of future events which is developed on the basis of past and current information[15]." Forecasting in the airline industry can be done on a macro-level, for example, the domestic US passenger traffic on an airline's route network in the upcoming year; the other extreme being that of forecasts made on a micro-level, for example, the number of Q-class passengers in a given O-D market for a given particular flight. Forecasting specifically applied to the airline revenue management problem entails obtaining an accurate estimate of the passenger demand based upon passenger trends on flights in the past and present. This demand can still be forecast on different levels; for example on a flight leg level, on a fare class level, or on an ODF level. In this thesis we will

[15] Pindyck & Rubenfeld, Chapter 8.

mainly concern ourselves with this third category--that of forecasting the expected number of passengers by leg or O-D market and fare class. Since there is no single best method for producing a "correct" forecast, various algorithms have been developed to predict the desired forecasted values (in our case, passenger demands) with good accuracy.

But one may ask where and why forecasting is necessary in the network airline revenue management problem. The methods used by the different seat optimization routines discussed below all require demand forecasts by fare class on future flights, in order to determine how many seats to protect for each fare class at every iteration of the optimization process (these forecasts may be on a flight leg basis or a path basis in each fare class, depending on the forecasting routine)[16]. Such a forecasting process is repeated at regular intervals during the booking process leading up to the flight[17]. As information on the current flight is obtained (i.e., actual bookings for the forecasted flight itself), the demand forecasts are revised to incorporate this new information. Thus, in order to reduce diversion and subsequent revenue dilution by protecting too many or too few seats for a given fare class, as accurate of a forecast as possible is desired at each instance where seat allocations are reoptimized.

Explained below are two of these different forecasting algorithms used in the PODS simulations; each one using different amounts of available information in various ways to produce passenger demand forecasts. It should be noted that the relative accuracy of the different forecasting methods available to us in the simulations is not being tested; instead, the relative merits of each method as applied to the airline network revenue management problem will be analyzed. Hence, we are not necessarily concerned with the precision with which our forecasts aim to predict the demand (although this is, in theory, very important in attempting to achieve the goal of revenue maximization), but rather with the effects on an airline's system revenue under different forecasting methods in a variety of different competitive scenarios. Furthermore, the compatibility of these forecasting methods with each of the different seat optimizers available to us in the simulation is also of interest.

[16] In the ensuing discussion, simply "flight" will be used to refer to a fare class on either a path or a leg; whether it is a path fare class or a leg fare class depends on the seat optimization routine implemented.
[17] In the PODS simulation, there are 16 such intervals, or *time frames*.

2.1.1 Pickup forecasting

Pickup forecasting is one level of detail higher than simple time series forecasting. A time series forecast in airline demand forecasting would simply be a weighted or unweighted mean of final departure bookings on a set of similar flights. The pickup forecaster goes one step further by including more information and averaging not just the final bookings for a flight, but actually the number of passengers picked up in the intervals preceding departure (i.e., the average incremental bookings received in each time interval before departure). The passenger demand forecast is made based upon complete booking information for previous flights, available in the historical database. This forecast could be made on a macro-level flight basis, but in our simulation it will be done on a more micro-level fare class basis, in which passenger demands are being disaggregated and forecasted by individual fare class (or virtual fare class, depending on which seat optimizer is being used--see Section 2.3 for discussions in this regard) and path or flight leg (depending on the seat optimization routine chosen). The reader should refer to Wickham[18] or Skwarek[19] for additional explanations describing the classical pickup forecaster.

To best illustrate the method of pickup forecasting, a small example is helpful. Let us examine the case that we are currently at some point in the booking process (here, at the beginning of Day 8, indicated by the dark line), for which we desire a forecast of demand for a particular fare class on a given flight on Day 9 (see Figure 2.1). Available to us from the airline database are the passenger bookings by booking interval for all flights which have already departed in addition to flights with incomplete booking profiles (these values are represented as the area to the left of the vertical line in Figure 2.1), although the pickup forecaster disregards the incomplete booking data from yet-to-depart flights[20]. For simplicity in explanation of this example, we will only consider the pickup in the last two booking intervals for each flight and the demand for only a single fare class, although these numbers can easily be extended to as many booking intervals and fare classes as desired in practical applications.

[18] Wickham (1995). Section 3.2 describes the different forecasting models tested, with Section 3.2.3 explicitly explaining the classical pickup forecaster.
[19] Skwarek (1996), Section 5.2.2.
[20] Other forecasting routines have been developed to take into account this information--see L'Heureux (1986).

The pickup model proceeds by taking simple numerical averages (which can be weighted if desired) of the incremental demands in the n^{th} booking interval before the flight departure for a certain number t of previous flights.

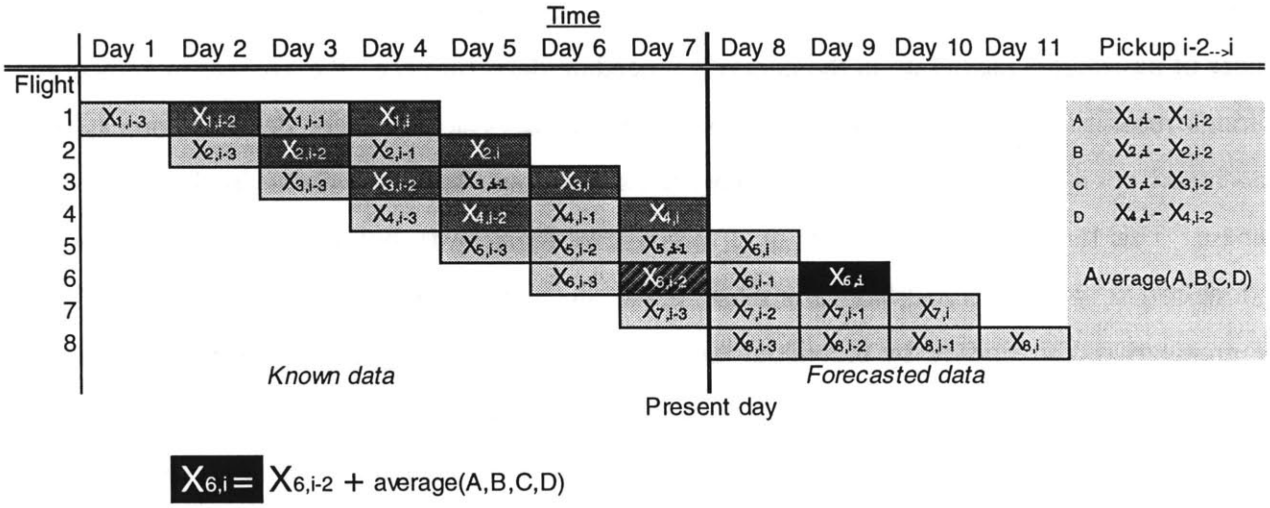

Figure 2.1: Pickup forecasting illustrated

Now, let us take the case where a forecast of demand at the end of the booking process is desired for Flight 6 (departing on Day 9). To determine this, the value for average pickup from time frame i-2 to time frame i is needed; it is obtained by averaging the pickup values corresponding to these time frames (i.e., the difference between the number of bookings in the medium-shaded blocks and the lightly-shaded blocks in Figure 2.1) for the t previous flights (four, in our example). Therefore, the values given in the blocks labeled A, B, C, and D are the pickup values for each flight, while the average is indicated in the same block for Flight 6. Now a final time frame forecast can be made for Flight 6 simply by adding this expected pickup average to the data value we already have for bookings already accrued as of time frame i-2 for Flight 6; thereby forecasting a demand of $X_{6,i}$ which is equal to $X_{6,i-2}$ plus the average pickup, as shown in the bottom of Figure 2.1.

In essence, the pickup model can be generalized and represented by Equation 2.1 (the notation corresponds to that used in Figure 2.1). Equation 2.1 is a generalized formula which can be used even if time-weighted averages are desired (e.g., more weight given to more recent flights); if all observations are to be weighted equally, then all values of W_j can be set to $W_j = 1$.

$$X_{M,i} = \left[\frac{1}{t} \sum_{j=1}^{t} W_j \cdot \left(X_{j,i} - X_{j,i-n} \right) \right] + X_{M,i-n} \tag{2.1}$$

where $X_{M,i}$ = total bookings after interval i for flight M

M = flight on which forecast is desired

n = number of time frames over which pickup is calculated

t = number of flights upon which forecast is based

j = flight index of flights upon which forecast is based

W_j = weighting value applied to individual flights (if desired)

The value of $X_{M,i-n}$ is just the number of forecasted bookings accumulated between the desired time frames i-n and i; it should not be confused with the number of cumulative bookings $X_{M,i}$ on the flight of interest (cumulative bookings can be found by adding $X_{M,i-n}$ to the actual bookings in the preceding interval). Usually, data storage by the airline will be of the cumulative bookings on a given flight, but pickup can easily be deduced from this information by simple subtraction. Therefore, once the forecasted values of $X_{M,i}$ are obtained by computing the pickup between the last time frame for which we have actual data and the desired departure day for each flight in which a forecast is needed, mean departure day demands are then available for each flight. These values can then be used by the seat optimization routine to make seat allocation decisions or bid price calculations.

It should be noted that several simplifications and assumptions have been made in the process of obtaining the forecast, most of which could be corrected for if the assumptions were relaxed. First, we are assuming that all flights with complete booking information (i.e., those flights which have already departed) behaved in the same manner and can be considered as "similar" to the flight of interest. In the airline world, this usually entails comparisons of flights occurring on the same day of the week, as demand by day tends to follow a pattern of non-uniformity (e.g., Mondays and Fridays generally experience higher demand than Wednesdays or Saturdays). Second, demand for air travel is highly subject to seasonal fluctuations, especially when this demand is disaggregated by fare class. Although an argument can be made for the fact that business travel has lower variance throughout the course of a year, empirical evidence for leisure travel definitely points to travel peaks and troughs (the peaks typically occurring during the summer months and the holiday season with the troughs occurring during the late fall and early spring). Therefore we must be selective in which historical information is to be used; taking into account too many prior flights may cross such seasonal boundaries, but the fewer data points upon which the forecast is based, the less reliable it will be.

To account for these factors, the PODS simulation was designed as a stationary process, with no flight trends by day of week or seasonality issues. Therefore, all flights can be considered to follow similar trends with respect to the assumptions outlined above. More specifically, the forecasting routine in the PODS simulation uses data from 26 previous flights in making forecasting calculations, judgmentally chosen to strike a good balance between the number of flights used for information (i.e., the larger the sample, the better the predictions) and how reflective they are of the current trends (i.e., too many flights in the real world would result in a forecast based on data that is not up-to-date with demand for the flight of interest).

Another forecasting method similar to the classical pickup model could alternatively be used; namely the advanced pickup model. This model combines not only information we have about flights which have departed (the medium-shaded blocks in Figure 2.1), but also about flights yet to depart which have some booking information available (i.e., any information available to the left of the line representing the present day in Figure 2.1. The origin of this advanced pickup model in terms of application to the airline revenue management problem can be traced back to L'Heureux at Canadian Pacific[21]; the method is also described by Skwarek[22].

[21] L'Heureux (1986) describes the theory behind both the classical pickup forecasting method and the advanced pickup forecaster, which builds upon the classical one by incorporating data from flights which have not yet departed.

[22] Skwarek (1995), Section 3.1.2.1 also describes and illustrates the data used by the advanced pickup model.

According to L'Heureux, there are both advantages and disadvantages of using such an advanced pickup forecaster over the classical pickup routine (where only flights with *complete* booking information are used as presented above). The classical forecaster as presented above has relative weaknesses in that it is influenced more by a single high-demand (or low-demand) flight during the booking period and it does not reflect seasonality trends well as data from more recent flights (the ones with incomplete booking profiles) is disregarded, whereas the advanced forecaster does a better job of accounting for this. On the other hand, it is less susceptible to periods of high or low booking activity. However, given that the trials in PODS are stationary processes, the classical pickup forecaster performs adequately, and was therefore used as the "base case" forecasting method choice.

2.1.2 Regression forecasting

In the previous section, a methodology for analyzing the expected demand by using historical information of flights which have already departed was presented. In this section, a different approach is taken in determining a forecast of the demand in a particular fare class; namely, regression forecasting. This forecasting approach is simply what its name indicates--a least-squares regression of demand by fare class for the flight of interest based on bookings received in that fare class as of some previous time frame. Further description of the regression forecaster is also given by Wickham[23].

Again, to illustrate the method, the same example case from Section 2.1.1 will be used. Let us assume we are still trying to obtain a demand forecast for the same flight (Flight 6) on Day 9. Again, only one fare class will be used along with only the two booking intervals prior to the flight for simplicity; both of these limitations could again be extended as needed. All of the same historical information available to us previously remains identical, and in fact, the same historical information will be used in this method as in the previous one. However, we are now interested in a simple linear regression model relating the final bookings with which a flight departed to the total accumulated bookings at some point in time prior to the day of departure. Figure 2.2 graphically illustrates this information.

[23] Wickham (1995), Section 3.2.2.

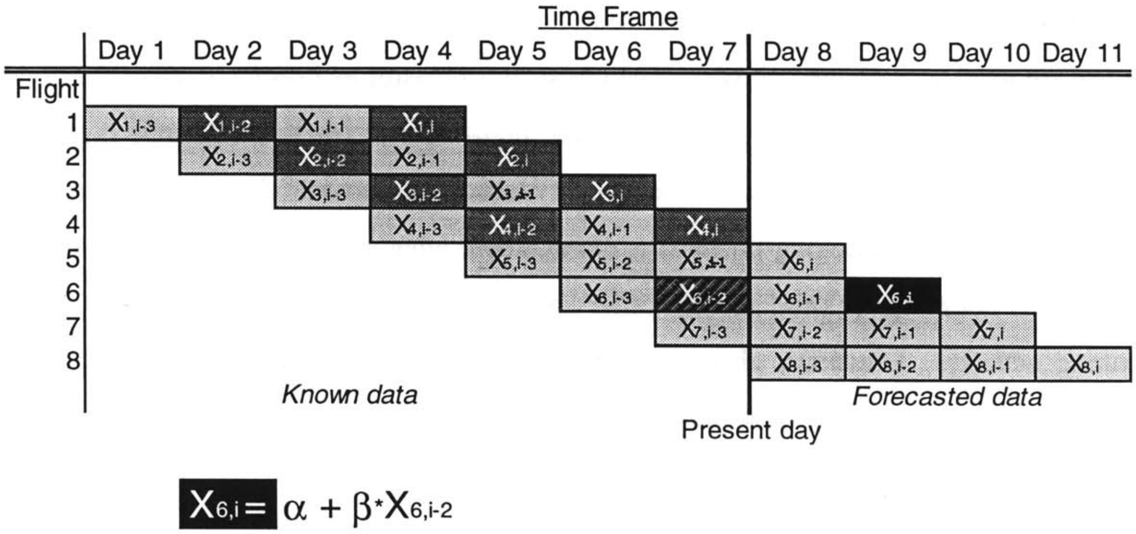

Figure 2.2: Regression forecasting illustrated

Here, a simple ordinary least-squares regression is performed using data from completed flights only, in order to obtain coefficients for the independent variables representing accumulated bookings at some specific time frame during the booking process. Therefore, in our example the medium-shaded blocks are the dependent variables (final time frame bookings for the flight), while the lightly-shaded blocks are the independent variables (bookings for the flight in some intermediate time frame). Once a linear regression is performed, coefficients relating these two variables are obtained; namely, α and β, which are just the slope and intercept. Once this equation is obtained, forecasts can be made for any flight on which we have completed booking information two intervals before departure. The regression method of forecasting can also be represented in mathematical form by either Equation 2.2 or Equation 2.3.

$$\mathbf{X}_{\mathbf{M,i}} = \alpha_{\mathbf{n}} + \beta_{\mathbf{n}} \cdot \mathbf{X}_{\mathbf{M,i-n}} + \varepsilon_{\mathbf{n}} \qquad (2.2)$$

$$\mathbf{X}_{\mathbf{M,i}} - \mathbf{X}_{\mathbf{M,i-n}} = \gamma_{\mathbf{n}} + \delta_{\mathbf{n}} \cdot \mathbf{X}_{\mathbf{M,i-n}} + \nu_{\mathbf{n}} \qquad (2.3)$$

where $X_{M,i}$ = total bookings after interval i for flight M

M = flight on which forecast is desired

n = number of time frames over which pickup is calculated

α_n, γ_n = coefficient on pickup variable for interval i

β_n, δ_n = intercept terms

ε_n, ν_n = error terms

Equation 2.2 states that total bookings on the day of departure for a fare class is just a function of the bookings received a particular time frame prior to departure. (Note that this can also be expressed as Equation 2.3, where the pickup value between these intervals, $X_{M,i}$ - $X_{M,i-n}$, is represented as a function of the bookings received a particular time frame prior to departure). Hence, ordinary least squares regression produces coefficients α_n and β_n; then substituting the known bookings in the previous time frame into the above equation produces a forecast of $X_{M,i}$. This differs somewhat from the pickup method in Section 2.1.1, but tends to be slightly more accurate because total bookings on day of departure is highly correlated with the bookings for that flight on some day before departure. In fact, it can be observed that pickup forecasting as illustrated in Section 2.1.1 is just a simplified form of regression forecasting with one less degree of freedom, in that the coefficient β is set equal to 1 and the coefficient α is equal to the average pickup for the prior flights being examined. However, despite the slight forecasting improvement in many instances, regression forecasting will be more suspect to error in flights with very high or very low demand.

2.2 Detruncation methods

Detruncation is a necessary procedure when we are trying to project demand estimates or demand distributions from historical data which has been restricted by some capacity constraint. Basically, in order for our seat optimization algorithms to operate correctly, we need to know the mean demand (and its variance) for a given fare class on a given path or flight leg (this again depends on which seat optimizer is used). However, the implementation of a capacity restriction on any given path or leg fare class does not allow us to directly infer passenger demand in cases where this capacity constraint is reached. That is, further passenger requests are being turned away--

passenger requests which should theoretically be considered in the demand estimate. Therefore, detruncation is a method by which an adjusted estimate is obtained by modifying the capacity-constrained data upon which the projection is based. In the airline revenue management case, detruncation lends itself to use on flights where the aircraft capacity becomes a constraining factor and thereby limits the demand accommodated by the flight or fare class.

As an example, let us assume we have an aircraft of capacity C; for which we know historically that some of the time the achieved demand is less than the aircraft capacity (bookings < C), while the rest of the time the demand for the flight exceeds this capacity (bookings = C, # refused requests ≥ 0). The former case poses no need for detruncation, as all potential demand is accommodated; as will be seen later, such observations form the basis for adjusting the capacity-constrained data. If the latter case is encountered, however, and the flight sells out because business travel demand was higher than expected and we did not protect enough seats for the high-fare passengers (i.e., there were leisure passengers who were taking some of these seats away), we have erred in our seat optimization and missed out on potential revenue by allowing low-fare passengers to take seats which could have been filled by higher-fare ones. (Note that if the capacity constraint was reached because leisure demand was too high, the results should not be dramatically affected, as we should still see the optimal number of seats protected for higher fare classes--there will just be a shortage of seats offered at the lowest fare class).

As can be seen, some method of adjusting our demand predictions to account for the capacity restriction is necessary. But upon what do we base the estimate of potential market demand? If it is based completely on all prior flights without somehow accounting for those in which the fare class of interest sold out, the predicted total market demand will always be underestimated to some degree, since we are neglecting to account for the refused requests in the demand projection. On the other hand, if we simply do not count those flights in which the fare class of interest sold out and use only flights with demand less than capacity to make our estimate, we will again underestimate the demand, as none of the high-demand observations will be used, and our estimates will only be based on the low-demand instances (additionally, the sample size from which we are drawing our data may become very small if the demand is generally high and closed fare classes are a common occurrence). Therefore, to avoid these pitfalls which are detrimental to the goal of revenue maximization, a detruncation method is used to adjust the demand estimate for the capacity constraint. All historical observances can then be used, provided that we correct for those observances (i.e., fare classes on flight legs or paths) constrained by capacity by detruncating (unconstraining) them to gain estimates of what the unconstrained demand would have been (i.e., demand with no capacity restriction). As in forecasting, there are numerous ways to go about this task. Discussed next are two of these algorithms used in the PODS simulation, both of which will be tested in Chapter 4.

Figure 2.3 below provides a good illustration of the need for detruncation in a case with capacity restriction. The distribution shown is of the number of passengers desiring a particular fare class on a given flight leg or path--this is assumed to be normally distributed. Given that we have no capacity restriction, the mean demand and its variance can be easily inferred from the graph--they are labeled as μ_U and σ.

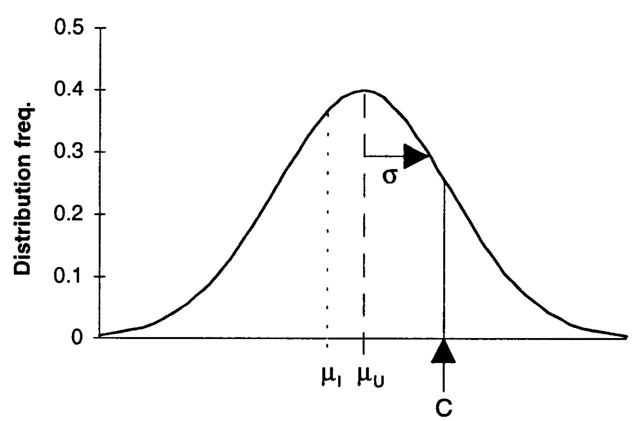

Figure 2.3: Demand distribution for all demand

However, a slight complication arises when a capacity restriction is placed on a fare class, as our seat optimization routines do with the fare classes they are optimizing. Also shown in Figure 2.3 is a capacity restriction C, represented by the solid vertical line. Observations to the right of this line are only recorded as values of C in the airline's historical database, although the true demand is somewhat higher. Therefore, the estimated average of our observations will come out lower, at μ_I, if we do not somehow account for this restriction. It should be noted that this fare class capacity restriction C is dynamically determined and can be adjusted in each booking interval, (i.e., each time the seat optimization routine is run), but regardless, some restriction is always chosen as a capacity limit.

As mentioned before, two cases can occur for a fare class on any given flight: (1) the realized demand does not reach the capacity constraint; and (2) the realized demand is equal to the capacity constraint. In the former case, we can directly infer the mean demand and its variance from the distribution as described earlier, after which the forecasting routine can then be implemented. It is in the latter case that direct inference of μ and σ cannot be done and detruncation must therefore be used. The reader is referred to Skwarek[24], in which he tested scenarios of an airline using no detruncation algorithm while one was used by the competitor; even at a demand factor of 0.9, the revenue impacts of the absence of such a system were greater than 3.5%. And when the demand factor was increased to 1.2, the impacts were altogether obvious, as more than a 50% revenue difference between the competitors was present. As stated, not detruncating the historical observations will cause low estimates of demand, and therefore not enough high fare class seats will be protected. To compound the effect, the competing airline who performs some sort of detruncation will protect more seats for the high-fare class passengers, thereby diverting the excess low-fare class demand to the competitor without the strict capacity limits on the lowest fare classes. Such results indicate the importance to the yield management system of the ability to adjust the demand predictions by unconstraining them for good revenue performance; Sections 2.2.1 and 2.2.2 discuss two of the different ways in which this is accounted for by the PODS system.

[24] Skwarek (1995), Sections 5.3.1 and 6.2.1. These simulations were done for the single-leg, two-competitor case; we would expect these effects to be even more pronounced for the network case of PODS 7b.

2.2.1 Booking curve detruncation

The first method to be discussed for detruncating the demand data and producing an adjusted forecast is booking curve detruncation (see Wickham[25] for detailed development of the method). Booking curve detruncation is a relatively straightforward process by which theoretical passenger booking curves are used to project demands of what the demand forecast would be if the observed data were not limited by the capacity constraint. Since the airline database contains booking information by fare class and time frame (as was the case in the forecasting discussions of Sections 2.1.1 and 2.1.2), a representative booking curve can be generated for each fare class and flight of interest by analyzing the passenger pickup in each interval.

The detruncation algorithm then proceeds by first estimating the booking curve for all unconstrained observations of fare classes on flights (i.e., all instances where the fare class did not close). The next step is then to calculate the ratio of bookings in each given booking interval i to the number of bookings in the preceding interval i-1 for a given fare class on a flight. This is done by averaging the total bookings in period i for all flights which did not close and dividing by the average of total bookings in period i-1 on the same flight, for all combinations of periods i and i-1 (see Equation 2.4 below). Doing this on a fare class and flight-specific basis for each unclosed flight will yield the average pickup ratios from interval i-1 to interval i (see Figure 2.4 for a graphic representation of these calculations). Equation 2.4 represents this pickup ratio for a particular fare class from booking interval i-1 to i.

$$\mathbf{R}_{i,i-1} = \frac{\sum_{j=1}^{M_u} \mathbf{X}_{j,i}}{\sum_{j=1}^{M_u} \mathbf{X}_{j,i-1}} \qquad (2.4)$$

[25] Wickham (1995), Section 4.2.3. The process of booking curve detruncation is described as a tool for unconstraining data on closed flights (note that the multiplier values used are the inverses of those presented here).

where $R_{i,i-1}$ = average booking curve multiplier from interval i-1 to i

$X_{j,i}$ = total bookings in interval i for flight j

M_u = flight index of unclosed flights

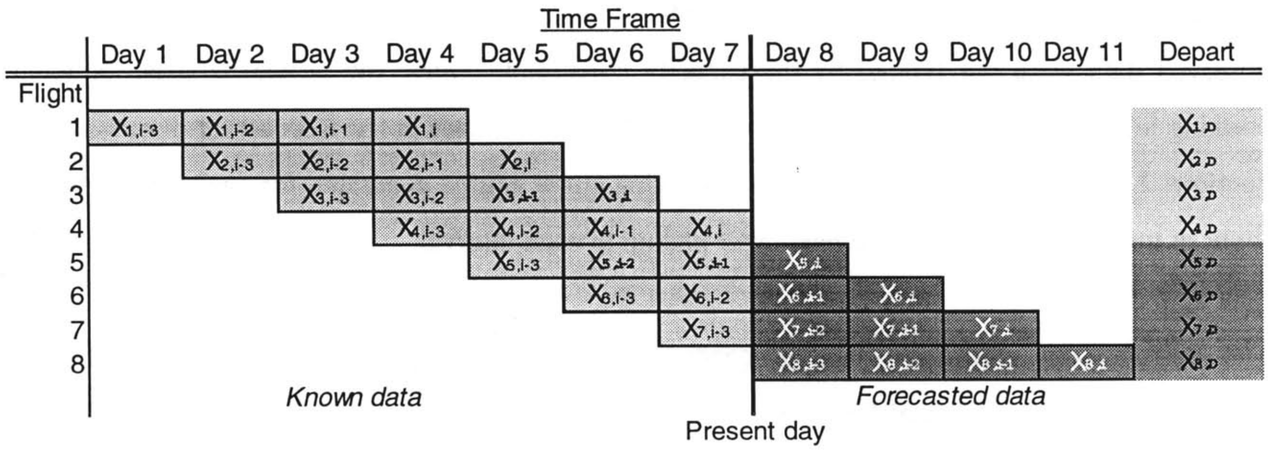

Figure 2.4: Booking curve multipliers

Once this series of ratios $R_{i,i-1}$ is known for each successive pair of booking intervals i and i-1, these values just need to be applied to the flight observations containing fare classes which closed at some point prior to departure. We know in which booking interval closure occurs for each of these flights, so the number of bookings in this interval (which is equal to the fare class capacity) just needs to be transformed by multiplying by the product of multipliers for all future intervals, obtained as values $R_{i,i-1}$ for each flight j from Equation 2.4. This transformation then yields the theoretical demand which would have occurred had no capacity restriction been in place on the fare class. So assume that for the flight of interest j, the fare class demand reached capacity in interval k. To obtain the unconstrained fare class demand for the departure, the transformation performed is represented by Equation 2.5 below.

$$\mathbf{X}_{\mathbf{M_c,D}} = (\mathbf{X}_{\mathbf{M_c,k}}) \cdot \prod_{\mathbf{i=k+1}}^{\mathbf{D}} \mathbf{R}_{\mathbf{i,i-1}} \tag{2.5}$$

where $R_{i,i-1}$ = booking curve multiplier from interval i-1 to i (calculated in Equation 2.4)

$X_{Mc,K}$= bookings on closed flight M_c at booking interval k (interval where closure occurred)

$X_{Mc,D}$= unconstrained departure day demand on closed flight M_c (had no closure occurred)

It should be noted that the product of the booking curve multipliers in Equation 2.5 is just the expected pickup ratio from the closure interval to departure. The above procedure can also be shown graphically, using the information from Equations 2.4 and 2.5 (see Figure 2.5 below). The average multipliers obtained from Equation 2.4 are shown in the first row; in the second row are the projected values by time frame for a flight j which closed in interval k; and in the third row are just the corresponding references to these values (i.e., their nomenclature from Figure 2.4).

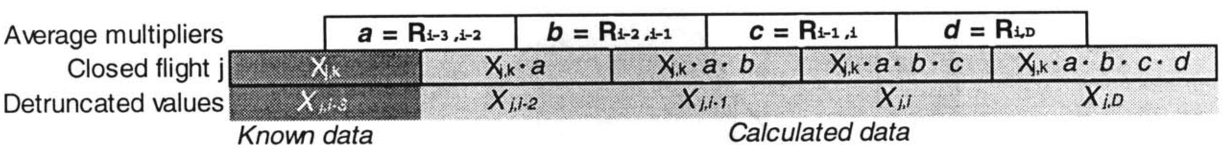

Figure 2.5: Booking curve multiplier transformations

Given this demand unconstraining procedure, we now have the mean demand values by fare class for flights on which the fare class capacity limitation was reached, as well as for those on which the fare class remained open. Based upon this information, better demand predictions can then be made by the forecasting algorithms in order to help the seat optimizer come up with more optimal seat allocations and capacity limits by fare class, thereby enhancing the revenue performance of the overall yield management system.

2.2.2 Projection detruncation

Booking curve detruncation does an adequate job of adjusting the desired data, but it relies on the need for tedious calculation of the ratio of pickup between every adjacent pair of booking intervals in order to compute the correct multipliers for the closed observations. Projection detruncation, developed by Hopperstad[26], is an alternative method which uses a more straightforward probabilistic approach to unconstraining the fare class demand. To begin with, assume we are at some point t in the booking process, and let us separate our historical flight and fare class observations as before into two groups--those which reached the capacity limit and those which did not. The mean μ and variance σ of the pickup from time t until departure are again simply calculated using the unclosed observations only. Again, we will assume a normal distribution for the booking requests on these flights, so with the calculated values of μ and σ, the distribution of the curve can be drawn. Given this distribution (which is assumed to be normal), we next draw a line corresponding to the capacity restriction; illustrated in Figure 2.6.

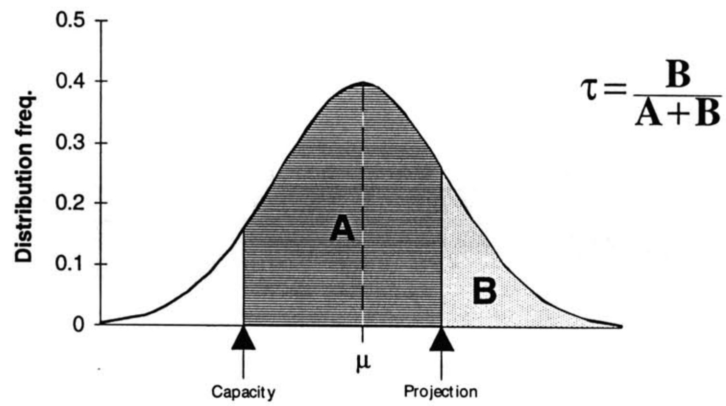

Figure 2.6: Constrained demand distribution for projection detruncation

The basis of projection detruncation relies on the fact that we are assuming that the conditional probability that we *underestimate* the unconstrained demand for a flight and fare class is a constant value (conditional on the fact that closure did occur for that particular flight and fare class). Simply put, this means that for the subset of observations in which fare class closure occurred (i.e., demand was larger than the capacity), there is a fixed percentage of those observations where the actual realized demand was, in fact, even higher than what we had been projecting. Fare class closure occurring puts us to the right of the capacity line in Figure 2.6, in shaded area A or B. Underestimating the demand for these observances corresponds to area B,

[26] Hopperstad (1997).

hence, the ratio τ is just the ratio of area B to the shaded areas A and B. Using probability theory, the value of τ can also be illustrated with Equation 2.6.

$$\tau = \frac{B}{A+B} \tag{2.6}$$

where τ = conditional probability that unconstrained demand was underestimated given that closure did occur

A+B = probability that demand > capacity (from Figure 2.6)

B = probability of underestimating unconstrained demand (from Figure 2.6)

Once this value of τ is determined and substituted into Equation 2.6, a unique solution of the quantity B can be found, since the area A+B becomes a known quantity when the capacity restriction of C is implemented. Finally, given that we now have a value for the area B (i.e., the probability of underestimating the unconstrained demand), our projected detruncation value of demand (i.e., μ and σ) can then be determined (also labeled in Figure 2.6). These newly obtained values for μ and σ are then used as input and the process is repeated until convergence of the μ and σ values below some specified tolerance occurs. Therefore, the projection scheme determines a probabilistic distribution of the demand, and then infers the unconstrained demand given the constrained demand and the conditional probability explained above. Similar to the case of booking curve detruncation, better demand predictions can then be made by the forecasting algorithms, as better estimates of the actual demand are available. Based on initial empirical evidence, a value of $\tau = 0.15$ was used, which provides a rather high adjustment to the projected values (i.e., 15% of the time the actual demand was higher than what was predicted). Several simulations were run to test the sensitivity of projection forecasting to the τ parameter, but it was found that under a wide variety of demand factors and yield management systems, changes in the value of τ had little effect on overall revenue performance (see Section 4.3 for more discussion).

2.3 Interaction with yield management methods

When attempting to optimize network revenues, a yield management system will first use a detruncation method to unconstrain (if necessary) any previously observed data as described above, from which the forecasting algorithm then uses these demand values as inputs to determine what the projected demand actually is in a given fare class on a given flight leg or flight path. From this, we have a good demand estimate for each fare class on a flight leg/path in a given O-D market, for which the seat optimizer can then be used to allocate the optimal number of seats by flight leg/path and fare class. This iterative process takes place at the beginning of each time frame during the booking process; it is therefore appropriate to consider the interaction of these forecasting and detruncation methods with the yield management systems. That is, how do the combinations of methods presented above interact with the seat optimizers themselves within each of the yield management methods? This is primarily the scope of Chapter 4. However, presented here is a brief discussion of the third integral part of the yield management system--the seat optimization; that is, how do we determine how many seats should be assigned to each of the different fare classes in the reservation system?

2.3.1 Brief descriptions of seat optimizers used

Five different seat optimizers will be tested in the PODS simulation, for which a brief description of each is provided here. The goal of this section is simply to give insight into the different methodologies by which these seat optimizers operate for an airline network, with Sections 2.3.2 and 2.3.3 discussing the interaction between them and the forecasting/detruncation methods described above. Detailed functionality descriptions of each seat optimization routine will not be provided; the reader is directed to Wei's thesis[27] in which a brief analysis of each of these various methods was performed.

Basically, the airline revenue management problem comes down to a simple linear program; maximize revenues (i.e., the product of passengers and fare for all O-D markets) subject to aircraft capacity and market demand constraints. This can be expressed in the form of Equations 2.7-2.9).

[27] Wei (1997). Chapter 2 gives detailed examples of each of the seat optimizers presented here.

$$\max \sum_{ij} \mathbf{P}_{ij} \cdot \mathbf{X}_{ij} \tag{2.7}$$

$$\text{s.t.} \quad X_{ij} \leq \text{demand}_{ij} \qquad \text{(for all O-D markets i-j)} \tag{2.8}$$

$$X_k \leq \text{capacity}_k \qquad \text{(for all flight legs k)} \tag{2.9}$$

where P_{ij} = fare offered in market i-j

X_{ij} = passengers carried in market i-j

X_k = passengers carried on leg k

Initial examination of this problem would tend to point to a network optimization scheme, where network revenues can be maximized (Equation 2.7) given the demand and capacity constraints from Equations 2.8 and 2.9. While theoretically correct, this proves to be very difficult to implement in real-world airline situations, for two primary reasons. First, the size of the problem becomes a constraining factor--for the network structure of a typical major US airline the solution time could be on the order of hours or even days. This does not lend itself to frequent reoptimization, something that would be necessary to continue to ensure revenue maximization during the booking process as actual demand is received. Second, data storage by the airlines has traditionally not been done on an ODF basis, but rather on a leg basis. Therefore, major implementation costs are required to change the system, or large-scale data transformation needs to be done to rectify the data; both of which require major investment. Therefore, the seat optimization schemes described in the following subsections have been developed to maximize the potential revenue while keeping implementation cost to a minimum (which does not necessarily correspond to full revenue maximization).

For a simple network example, let us examine the (oversimplified) situation in Figure 2.7, where there are three cities (A, B, and H), with one of them (city H) serving as a connecting hub.

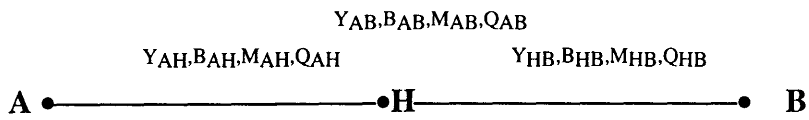

Figure 2.7: Graphical depiction of simplified network

The airline will operate service as one flight over each of the two legs (A-H and H-B), while there are three distinct O-D markets for passenger travel (A-H, H-B, A-B), along with four separate fare classes offered (Y,B,M,Q) in each of the three O-D markets. The highest (i.e., most expensive) fare class is Y, while Q is the deep discount coach fare. In essence, the optimization of the linear program above comes down to determining how many seats should be reserved for each fare class on each leg or path based upon our projections of ODF demand, with nested protection limits (i.e., seats reserved for a fare class available to that fare class as well as any classes above it) thereby determined for each fare class.

Seat Optimizer	Estimation	Control
EMSRb	leg	leg
VEMSRb	leg	leg
HBP	leg	leg/path
DAVN	path	leg
Netbid	path	path

Table 2.1: Demand estimation/control[28]

Demand estimation and control can either be performed on a leg basis or a path basis depending on the choice of seat optimization routine; Table 2.1 above shows how each seat optimization routine discussed below estimates demand and controls seat allocation (on either a leg or path basis). The methodologies for doing so are numerous; several of them will be explained in the subsections below.

[28] Hopperstad (1997).

2.3.1.1 *Expected marginal seat revenue (EMSRb)*

The EMSRb (expected marginal seat revenue) yield management system is the first seat optimizer presented here; it is based upon the original EMSR model developed by Belobaba[29]; this original EMSR model was later modified into the EMSRb model by Belobaba[30]. As with all seat optimizers outlined in this thesis, the basic goal is to determine the optimal fare class protection limits (i.e., how many seats to reserve for each higher fare class) or the optimal bid prices (i.e., the cutoff price for which passengers are accepted if their fares exceed it) in order to maximize system revenues. Using EMSRb, the first step is to determine the probability that the next passenger booking on a leg (e.g., A-H or H-B in Figure 2.7) will occur in a given fare class. This is done using probability theory and the known quantities of mean demand and its variance by fare class, based on historical observations and our generated forecasts. Once we know this probability for all fare classes, a corresponding expected revenue value of filling the seat with this passenger (given our pricing structure) can be determined. We then protect a seat in the highest fare class for which this expected marginal seat revenue is greater than the fare level in the next lowest fare class; the top-down aggregation of this process will thereby produce nested protection limits for each fare class.

Fare class	Market A-H Fare value	Market H-B Fare value	Market A-B Fare value
Y	$700	$800	$1,200
B	$550	$600	$1,000
M	$400	$400	$800
Q	$200	$250	$400

Table 2.2: Hypothetical fare class table

The EMSRb yield management algorithm operates on a flight leg basis; that is, each leg is optimized individually, based on demand distributions for the flight leg itself (irrespective of passenger O-D markets). In the simple network case for only a single connecting O-D market

[29] Belobaba (1987), Chapter 5 describes the mathematical foundations of the EMSR model, while Chapter 7 analyzes the testing and simulation of the model.
[30] Belobaba (1992). A new EMSR heuristic was added to compute the *joint* protection levels for the higher fare classes in the EMSRb optimizer, as the simple EMSR optimizer's nested booking limits could be sub-optimal for lower fare classes as only one demand density at a time was being considered.

(i.e., which contains two distinct local O-D markets in addition to a connecting one as in Figure 2.7), this is done by considering each leg fare class individually in the calculations, with demand-weighted averages of the local and connecting traffic used to compute the adjusted fares in each fare class. Following this, the optimal booking limits by fare class for each leg are then computed. Hence, for travel in a connecting market, seats must be available in the given fare class on *both* legs of the flight. What this method fails to account for is the fare differences which can occur in different O-D markets when a network is being optimized. Looking at the hypothetical fare table in Table 2.2, a local market whose Y-class fare is $700 but whose Y-class is closed is turning away potential Y-class customers on the long-haul connecting flight whose Y-class fare may be $1200. Additionally, a local market Q-class customer may pay $200 and thereby cause a $400 Q-fare connecting customer to be spilled. In essence, more seats may possibly be protected for the lower-fare local market demand which does not end up materializing if demand is low on some legs, while long-haul connection passengers are simultaneously turned away. One attempt to correct this problem is addressed by the next seat optimization routine presented below.

2.3.1.2 "Greedy" virtual nesting (VEMSRb)

The VEMSRb (virtual class expected marginal seat revenue) yield management system uses the same basis as the EMSRb system for calculating the booking class seat protection values; however, it uses "virtual" fare classes for the seat assignment process as a way of adding a degree of freedom to try to correct the problem presented at the end of Section 2.3.1.1. Originally developed at American Airlines[31], the optimization is carried out in identical fashion as for EMSRb; only here a set of virtual fare class "buckets" is judgmentally chosen and used, based on the magnitude of the fare in a given class. A possible example of virtual fare class divisions from the hypothetical example in Section 2.3.1.1 is given in Table 2.3.

[31] Smith et al. (1992). Virtual nesting, or the process of clustering ODFs into groups of virtual buckets, is described with a small example of how it was implemented at American Airlines.

Virtual fare class	Fare value	Mapped ODFs
Y1	\$1,000+	Y_{AB}, B_{AB}
Y2	\$800+	Y_{HB}, M_{AB}
Y3	\$600+	Y_{AH}, B_{HB}
Y4	\$400+	B_{AH}, M_{AH}, M_{HB}, Q_{AB}
Y5	\$200+	Q_{AH}, Q_{HB}

Table 2.3: Fare class table

Therefore, individual leg optimizations are not done strictly by existing fare class as in EMSRb, but rather by each leg's fare class "mapping" to some virtual fare class based on the relative magnitude of its fare level (there may be up to as many virtual buckets as ODFs on all flight legs). The advantage of using virtual bucketing is that hundreds of ODFs can be aggregated into a more manageable number of buckets. Also, doing this thereby favors the more expensive, long-haul flights, since a decidedly lower fare class on a connecting flight may be equivalent to the top fare class of one of the local O-D markets (e.g., the \$800 M-class fare in the connecting market A-B is mapped to the same virtual bucket as the \$800 Y-class fare in local market H-B). Similar to before, for travel in a connecting market, seats must be available in the given virtual bucket on *both* legs of the flight. One problem with such an algorithm (that favors the longer-haul flights) is that if demand is high on various legs, it could be more profitable to carry two local passengers rather than a single connecting passenger[32]. In essence, no passenger displacement costs are calculated; a modification for this is attempted in the following three seat optimization routines.

2.3.1.3 Heuristic bid price (HBP)

The heuristic bid price (HBP) yield management system, like the EMSRb and VEMSRb ones, also operates on a leg basis; in fact, the theoretical protection limit calculations are still performed identically as before. However, instead of calculating seat protection limits by leg/path and fare class (as done in EMSRb, VEMSRb, and DAVN), this approach operates on the principle of calculating bid prices and comparing the passenger fares in the decision as to whether to accept or

[32] Belobaba and Hopperstad (1997). VEMSRb methodology was shown to perform inadequately (revenues decreased by more than 0.8%) at a high demand factor (DF 1.2), even when the competition was using a simple fare class (i.e., EMSRb) approach. This stems from the fact that VEMSRb favors connecting passengers at high demand factors, thereby spilling high yield local passengers to the competition.

reject their request. This method was developed by Belobaba[33]; in it a bid price is calculated for each leg in order to determine the network contribution of the passenger's origin-destination combination by accounting for their downline displacement cost. The bid price is equal to the expected marginal seat revenue (EMSR)[34] of the last seat in inventory on that leg for the local passenger, while the bid price for the connecting passenger is this value plus the product of the percentages of local passengers on both connecting legs multiplied by the EMSR of the other leg (for the two-leg case, these values are given by Equations 2.10 and 2.11).

$$\textit{Leg 1}: \ BP_1 = EMSR_1 + d \cdot EMSR_2 \tag{2.10}$$

$$\textit{Leg 2}: \ BP_2 = EMSR_2 + d \cdot EMSR_1 \tag{2.11}$$

where BP_i = calculated connecting passenger bid price for leg i

$EMSR_i$ = expected marginal seat revenue on leg i

d = product of percentages of local passengers on *both* legs[35]

To determine whether or not to accept passenger requests, a simple comparison is performed and passengers are accepted if the fare they desire is greater than the bid price on their leg (for local passengers), or if the desired fare is greater than each individual leg's calculated bid price for all traversed legs (for connecting passengers). This method is advantageous in that it does not require network optimization but does account for downline displacement impact of connecting passengers; so incremental revenue can be gained with minimal implementation cost[36]. Additionally, fewer local passengers tend to be rejected than in the "greedy" case of VEMSRb, as the bid prices account for passenger displacement costs. However, one main disadvantage to this and any bid price control scheme is that of the lack of seat protection control. Because passenger

---

[33] Belobaba (1998). Instead of simply comparing the EMSRs on a leg or combination of legs to determine seat protection, the leg-based heuristic approach compares the EMSR from a passenger less the cost of the possible displacement of other passengers by that passenger.

[34] See Belobaba (1987), Belobaba (1992), or Wei (1997) for calculation of EMSR.

[35] The value d can also be thought of in another way; namely, as the probability of displacing a local passenger on each of the two flight legs. If we are assuming a 50-50 split between local and connection passengers on each flight, the value of d is simply $0.50 \times 0.50 = 0.25$.

[36] Belobaba (1994). Dynamic virtual bucketing, static virtual nesting, and stratified buckets can all be applied to the EMSR heuristic, although in PODS only static virtual nesting is tested, as the implementation cost of the other methods is substantially higher.

acceptance/rejection is simply compared with a fare, there is no limit on the number of passengers which may be accepted between recalculations of the bid prices. As for the HBP seat optimizer, Belobaba estimates that the gains achieved by such heuristic O-D control amount to nearly half of the possible revenue gains from the best network O-D control[37]. Hence, the next two methods presented aim for improvement upon the HBP results by using a network optimization scheme.

2.3.1.4 Displacement adjusted virtual nesting (DAVN)

The displacement adjusted virtual nesting (DAVN) yield management system operates on the principle of solving the revenue maximization linear program (illustrated in Section 2.3.1) to determine the shadow prices for each leg in the network (shadow prices are defined as the expected revenue increase from relaxation of the capacity restriction by one unit). Next, pseudo fares which take into account passenger displacement costs are computed for each leg. For local passengers the pseudo fare is just their desired fare on the leg, while for connecting passengers, the pseudo fare for a given leg is the total fare less the shadow price on the other leg traversed (see Equations 2.12-2.15 below)[38]. Hence, for two connecting legs i and j, we have the following equations.

<u>Leg i:</u>	*Local pax*:	PF^i_L = Fare on leg i	(2.12)
	Connecting pax:	PF^i_C = (Fare on leg i) - (Shadow price on leg j)	(2.13)
<u>Leg j:</u>	*Local pax*:	PF^j_L = Fare on leg j	(2.14)
	Connecting pax:	PF^j_C = (Fare on leg j) - (Shadow price on leg i)	(2.15)

where $PF^{(i/j)}_{(L/C)}$ = pseudo fare for local/connecting passenger on leg i/j

Once these pseudo fares are known, virtual buckets are used as in VEMSRb, but now they are chosen and reallocated based on the calculated pseudo fares (as opposed to the actual fares),

[37] Belobaba (1994).

[38] These pseudo fares can be reoptimized as often as desired, although the default used in the PODS simulations is to perform reoptimization only once at the start of the booking process.

thereby allowing simulation of the network displacement costs[39]. This type of seat optimization procedure has been shown to perform very well (see Belobaba and Hopperstad[40] and Wei[41]), and it is beneficial in that it also accounts for the connection passenger displacement costs on a network basis (as opposed to by flight legs)[42]. However, performing network optimization is extremely time-consuming, especially with the number of cities in real airline networks, and good forecasting on an ODF basis is much more difficult, as forecasted demand averages tend to be very small with large variance for any particular ODF (see Williamson[43] for discussion).

2.3.1.5 Network bid price (Netbid)

Similar to the displacement adjusted virtual nesting procedure described in the previous section, the network bid price (Netbid) seat optimization algorithm is another optimization routine which keeps historical data and predicts demand forecasts on an ODF basis. However, it uses a bid price approach to compare passenger fares as did the HBP method (Section 2.3.1.3), rather than calculate seat protection limits as DAVN does. It also solves the network linear program (Section 2.3.1) as before, to determine each leg's shadow price[44]. However, in this case, passenger acceptance/rejection is determined by a simple comparison of the passenger's desired fare and the corresponding leg shadow prices. Local passengers should be accepted if their fare is greater than the shadow price on that leg, while connecting passengers should be accepted if their fare is greater than the sum of the shadow prices on all legs traversed. Although this is a relatively simple optimization scheme to implement and it performs well in many cases, it also encounters the same time and data complexities that DAVN does; namely, the difficulty of making forecasts on an ODF basis and the difficulty of performing frequent network optimization (the bid prices

---

[39] The virtual buckets can be disaggregated as discretely as having leg-specific virtual buckets, although the default used in this thesis is to have a single network-wide set of virtual buckets based on the calculated pseudo fares.
[40] Belobaba and Hopperstad (1997). DAVN was shown to increase revenues by more than 1.75% at DF 1.0 against a leg fare class (EMSRb) seat optimizer used by the competition; and even a 0.5% increase was experienced at DF 1.0 when both competitors chose DAVN as the seat optimization routine.
[41] Wei (1997), Section 4.3.4. The DAVN seat optimizer in the simulations run provided revenue increases on the order of 0.8%-1.3% under the medium and high demand factor cases (DF 1.0 and 1.2, respectively).
[42] It should be noted that although DAVN performs seat inventory control on a leg basis, the forecasts are obtained by path.
[43] Williamson (1992), Section 4.2.2. This is the "small numbers" problem inherent in ODF forecasting.
[44] See Simpson (1989) for a small example where the network shadow prices are determined.

must be updated frequently[45] to avoid revenue dilution since no booking limits are in place). The reader can refer to Williamson[46] or Swan[47] for a more detail on the network bid price approach.

2.3.2 Forecasting and detruncation for yield management methods

The seat optimization routines discussed above are implemented in the overall yield management process once we are at the stage where booking limits (i.e., seat protection values) or bid prices are needed for each fare class on each path/leg in the airline network. Once a specified seat optimization routine has been chosen, independent choices are made with regard to the forecasting and detruncation methods by which the demand projections will be made and upon which the seat optimizer will base its calculations.

The detruncation routine answers the question of how to account for closed flights/paths/fare classes in estimating demand, while the forecasting algorithm attempts to predict how many passengers are expected in each path/fare class. In the forecasting process, there are two time variables; the time of the booking and the time of consumption (i.e., day of flight)[48]. This is what makes airline revenue management forecasting an intricate task--the consumption of air travel does not take place immediately following purchase. Hence, for a given day of consumption (i.e., flight) predictions of the forecasted demands by fare class must be made. Finally, the seat optimization algorithm then decides how many seats should then be protected for each path/fare class, based on the other information. The sequence of these algorithms is illustrated in Figure 2.8.

[45] Reoptimizations after every 10 bookings will be used in the PODS simulations tested, although we would ideally like to reoptimize the bid prices after every booking.
[46] Williamson (1992). Section 4.2.3 provides a more detailed description of the deterministic network bid price application.
[47] Swan (1994). In this presentation, the origins of bid pricing are presented, along with a small example of its methodology.
[48] As eluded to by Wickham in Section 1.5.1.

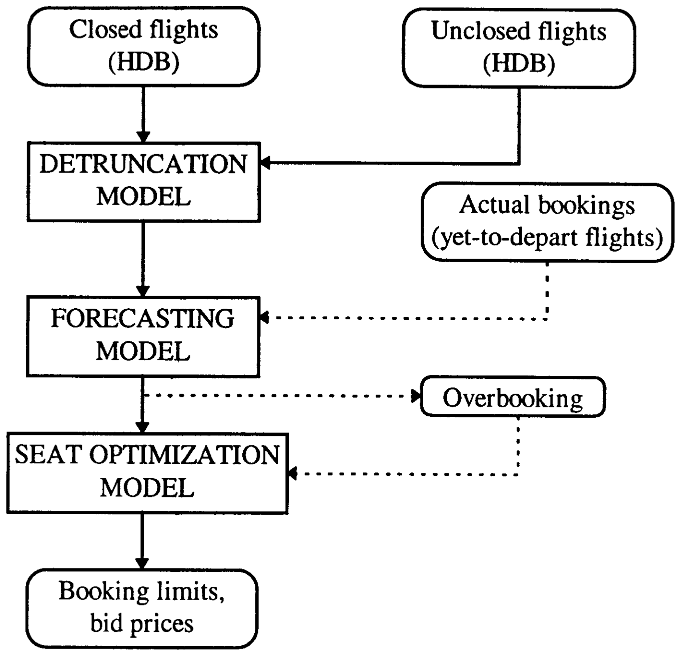

Figure 2.8: Sequence of model interaction

The process illustrated in Figure 2.8 above is followed immediately after the beginning of each time frame (or reoptimization). Both closed and unclosed flight information from the HDB are used as inputs to the detruncation model. In essence, the trends observed on the unclosed flights are used to project similar trends for the intervals following closure on flights in which closure occurred at some point prior to departure. Once the demand has been detruncated for all closed flights, the forecasting routine then incorporates this historical information along with the current booking information for the flight being forecast to project a departure day estimate of demand. (It should be noted that neither of the two forecasting algorithms tested in this thesis use information from other flights which have not yet departed, although the advanced pickup forecaster alluded to does). Next, if desired, an overbooking algorithm is applied to the projected departure day estimates of demand to account for the expected rate of no-shows (since no no-shows are being used in the PODS simulations, no overbooking need be done). Finally, these estimates of demand for each fare class on the specified flight are given as input to the seat optimizer, from which seat allocations or bid prices by fare class are computed. The next chapter details the background of the PODS simulator in which these algorithms are used.

Chapter 3

The PODS Model

PODS is an acronym for Passenger Origin-Destination Simulator, a complex passenger choice and yield management competitive simulation model developed at Boeing[49]. It evolved from the Boeing Decision Window Model[50], another simulation developed at Boeing to model passenger choice in an airline market, given schedules, airline characteristics, and a variety of other factors. While the Decision Window Model models passenger preferences for flights based on frequency and airline "image" disparities, it omits two important variables; namely, the fare (or set of fares) offered in the market and the capacity restrictions for the aircraft (or on the fare classes themselves). The PODS model has these added enhancements built in, and while an integral part of the simulation replicates the passenger choice routine of the Decision Window Model, PODS is also is capable of simulating competitive yield management practiced by one or more of the hypothetical airline competitors.

When PODS was originally developed, it was able to simulate a single flight leg (i.e., one O-D market), from which the competitive impacts of yield management implementation could be analyzed. Since then it has been expanded to simulate a typical airline network of spoke cities interconnected by hub airports, in which the hypothetical competitors have not only a wide variety of choices for their yield management system but also one in which lower-level inputs such as forecasting methodologies can be altered. Furthermore, the network structure of PODS allows for passenger choice among paths and fare classes; in essence, the ODF demands are not independent but rather interrelated. The ensuing discussion will therefore present a general overview of the current PODS model, the process by which the simulation is performed, along with some of the input parameters for the methodology (the reader can again be referred to

[49] See Hopperstad (1996) for a full-scale description of PODS version 6.
[50] See Boeing's "Decision Window Path Preference Methodology Time Mode" description.

Wilson[51] or Skwarek[52] for more detailed discussions about the intricacies of the simulator itself).

3.1 A brief description of PODS

The primary goal of the PODS model in this research is to give the user the ability to analyze competitive effects of yield management for an airline network, accomplished by running the PODS simulation model with a specific set of input parameters. Within any of these individual simulations, specific nomenclature is adopted and will be explained here. A basic schematic of the general PODS architecture is presented below in Figure 3.1.

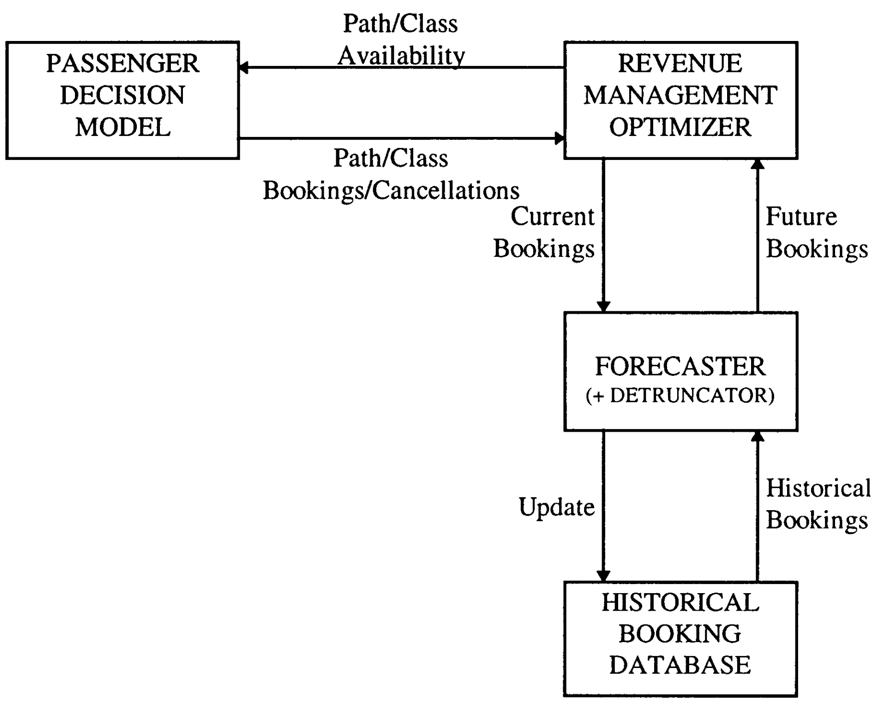

Figure 3.1: PODS basic schematic diagram[53]

In a macro-level sense, PODS is composed of four main components interconnected by three main feedback loops, as illustrated in Figure 3.1. First, there is the interaction between the

[51] Wilson (1995). In Chapter 3, Wilson gives detailed descriptions of the PODS system architecture, demand generation, and passenger assignment.
[52] Skwarek (1997). Skwarek also provides similar descriptions in Chapter 4, but with a heavier emphasis on its relevance to forecasting and detruncation aspects.
[53] Hopperstad (1997).

passenger decision model and the revenue management optimizer (seat optimizer), in which path and class availability are used as inputs in determining passenger decisions within the passenger decision model, while passenger assignments and cancellations are fed into the seat optimizer to be used during the subsequent optimization of seat allocations. Meanwhile, the seat optimizer has interaction with the forecaster (and detruncator); current bookings obtained from passenger assignment/cancellation are input into the forecasting routine to aid the estimation of future demand, while expected future bookings are given as output from the forecasting routine to the seat optimizer so that better seat protection or bid price decisions can be made. Finally, the interaction between the forecasting routine and the historical database is evident in that once the current bookings are fed into the forecaster, they are also used to update the historical database with the current information. At the same time, it is also these values of historical bookings that are used by the forecasting routine to project demand estimates. This macro-level interaction occurs in a systematic manner for each PODS case and trial; this will be described next (refer to Figure 3.2).

Simulations in PODS are run by *cases*; within each *case*, a specific combination of forecasting, detruncation, and seat optimization routine is specified for each competing airline (other parameters are also included--see Tables 3.2-3.4 in Section 3.3 for a complete listing). A *case* therefore contains the set of parameter inputs being tested under the simulation, the final output of which will provide the user with an idea of the revenue and passenger load impacts under these input conditions. Each *case* consists of 20 distinct *trials*; within any particular *trial*, there are 600 *samples*, where a given *sample* corresponds to one set of flight departures (since each market only contains one departure or connection possibility per day, one *sample* also corresponds to one day). We see therefore that a total of 12,000 samples are run, as we have 20 separate *trials* of 600 *samples* each. This disaggregation is done simply to reduce (by virtually eliminating) the correlation among the samples, as each sample is dependent on the existing conditions of the one immediately preceding. Running all 12,000 samples together would cause even the final one to be affected (albeit minimally) by the first, while separating them into distinct groups permits us to overcome this correlation problem. The grouping choice of 600 is used so that enough samples will still be run to have an outcome with a sufficient number of steady-state data points for statistical comparison[54].

During the simulation process, the first 200 of the 600 observations (flights) within each trial are discarded in order to remove any initial condition effects, since the ODF demands for the historical booking database within each airline's yield management system are estimated and input at the beginning of each *trial* and will therefore not have taken into account any previous flight information at that point. (The initial database is just a collection of best "guesses" for flight forecasts based on the inputted passenger booking curves, but these "guesses" are not based on any complete or known flight information). Hence, the accuracy of the initial *samples* is subject to high variance and randomness, since no steady-state period will have been reached (i.e., the period where the initial conditions have no effect). However, as time goes on and more samples are accumulated, actual flight information is entered into the database--this information then serves as historical data upon which more accurate flight forecasts for the next sample can be made. Results from earlier simulations have indicated that 200 discarded trials is sufficient to allow the seat optimization and forecasting routines to equilibrate so that the initial conditions have negligible effect.

In order to run the yield management simulations, realistic passenger booking data and airline network data are necessary. PODS generates this passenger demand using Boeing's Decision Window Model concept[55]. In generating the passenger demand, the Decision Window Model performs three crucial steps. First, a decision window is modeled--this is a time window within which the passenger is willing to travel. Second, given this time window, the possible paths for the passenger are generated, where a path is the flight or flight sequence that will take the passenger from their origin to their destination. Finally, given the feasible paths which correspond to the passenger's decision window, the first-choice path preference is systematically generated, using factors such as airline image and path quality (i.e., number of stops, connections, etc.). This first-choice path is based on a probabilistic calculation of each of the feasible paths from the third

[54] Lee (1998). Section 3.1.2 illustrates several cases which were tested; the best one was chosen to be that in which the revenue disparity was the smallest between the two airlines when they each used identical input algorithms.
[55] Refer to the Boeing Decision Window synopsis for the factors which are considered in each of the three following steps and the internal calculation mechanism.

step. Once these passenger demands are established by O-D market, stochastic booking processes are followed for each passenger type to simulate actual demand; while on the supply side, seat allocations are optimized to generate the maximum possible revenue from the expected incoming demand. As described above, PODS runs simulation *cases* in a series of 20 *trials*--for any single PODS *trial* (i.e., sequence of flight departures for each leg flown), a specific process by which passenger demands are generated and booking limits are set is followed by the simulation routine for each departing flight, graphically illustrated by Figure 3.2 below.

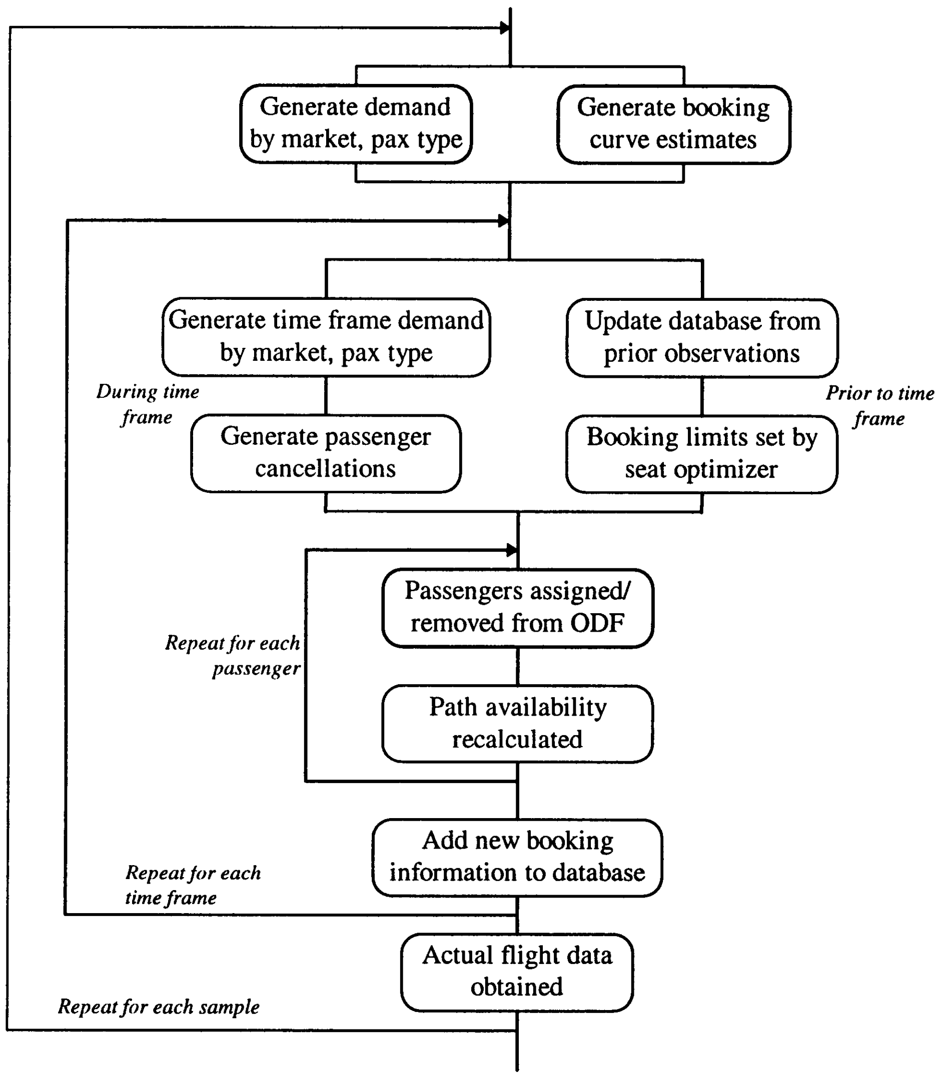

Figure 3.2: PODS flow chart for a single trial[56]

[56] Adapted from Hopperstad (1996).

To begin with, an initial historical demand database is constructed, since no previous flight information exists at the beginning of a trial from which one can be generated. This is done by making initial demand projections by market and passenger type based on the booking curves initially given as input to the system. The simulation then proceeds for each *sample* of interest, simultaneously performing both supply and demand functions; on the supply side, demands are forecasted and seats allocated accordingly by fare class; and on the demand side, passenger loads (i.e., market demands) based upon simulated passenger choice are generated for each flight as time passes.

For each flight of interest within a *trial* there are a number of *time frames* during which the PODS simulation performs various functions. First, the seat optimization routine is invoked at the beginning of each time frame to determine the seat protection limits and/or bid prices. This is done in two main steps--by making demand projections using the forecasting (and, if necessary, detruncation) algorithms chosen for the particular case (these forecasts are based on the most recent data available in the historical database), and then by setting protection limits by path/leg and fare class based on these forecasted values. Next, passenger bookings/cancellations are spread randomly throughout the entire time frame; passengers are accepted or rejected from a given ODF based on passenger choice as discussed above. Finally, the historical database is updated with these newly acquired passenger values at the end of each time frame, after which the simulation process then loops back to the beginning of the time frame loop for the next subsequent time frame (as illustrated in Figure 3.2).

Once these processes are completed for all time frames on all paths or flight legs (depending on the seat optimization routine used), the supply-demand routine is exited and actual flight data is available; repetition of this process for each *sample* produces data from which total system revenues and loads by leg or market are computed for the entire *trial*. This process is then repeated for each new *trial*, and upon completion of the 20 *trials*, simple numerical averages are calculated for the system revenues and loads--these are averages of the 8,000 *samples* used (20 trials $\times$ 400 samples per trial) and are the revenue and passenger load values used in the analyses in Chapter 4.

3.2 Airline network used in simulations

Initial tests using the PODS model[57] were originally done by simulating a single-route case of two competing airlines serving one origin-destination city pair, where leg-based seat optimization methods could be used and passenger choice effects could be examined and tested. The benefits of implementing a yield management seat inventory control mechanism on such a configuration (as opposed to not using one) were shown by Wilson[58]; while Skwarek[59] analyzed and detailed the additional effects of forecasting and/or detruncation method changes under a similar one-leg configuration (although simultaneous seat optimization method changes were not analyzed). Since then, however, the simulation capabilities of PODS have been expanded in order to more accurately model a real-world competitive scenario. Lee[60] has used this new setup to extend Wilson's study of the benefits of yield management to such a network case, modeling several competitive seat optimization changes by airlines on a network basis. This thesis will also use a similar network scenario which more accurately models an airline's entire network system to gain a better understanding of the *contribution* of forecasting or detruncation to the revenue increases which were obtained under simulations using different seat optimization algorithms, as well as the resultant consequences of the *interaction* of forecasting and detruncation routines with the different seat optimization methods.

The network layout of the updated version of PODS used in this thesis is described here, and can be referenced in Figure 3.3 below. The network itself is composed of six spoke cities, in addition to two airport hub cities; each hub serving as the focal point for connections of flight legs for one of two competing carriers. Such a setup models the airline trends which currently exist in the domestic US, where a variety of routings in any given O-D market pair are possible through the hub airport of any of a number of carriers. A good example occurring in the domestic US is evident on transcontinental travel--a passenger originating in Boston or New York bound for San

---

[57] PODS version 7b is used for all simulations in this thesis; older versions were used in previous research.
[58] Wilson (1995).
[59] Skwarek (1997).
[60] Lee (1998). Section 4.1 describes the network O-D based control tested, as opposed to the leg-based control done in the Wilson study.

Diego or San Francisco will have many path choices, depending on which connection hub will be traversed en route (this choice is a function of the airline chosen). Among the six spoke cities and two hubs in the simulated network, 54 different O-D markets can be constructed, interconnected by 24 separate flight legs (12 on each airline)--any combination of two of the eight cities (i.e., six spokes, two hubs) is possible as an O-D city pair with the exception of travel from one hub to the other (in the simulation, neither airline provides service to the competing airline's hub city).

On the demand side, two passenger types are considered--business and leisure. This is in contrast to disaggregating demand by individual fare class in a market, which would require the (more unrealistic) assumption of independent demands by fare class. This is one aspect which distinguishes the PODS simulation from other traditional models; while other simulations tend to assume independent demands by ODF, the PODS formulation generates demand correlated across passenger types and markets. Hence, by dividing the demand into business and leisure categories we can better model the actual airline world in addition to allowing for the possibility of sell-up[61].

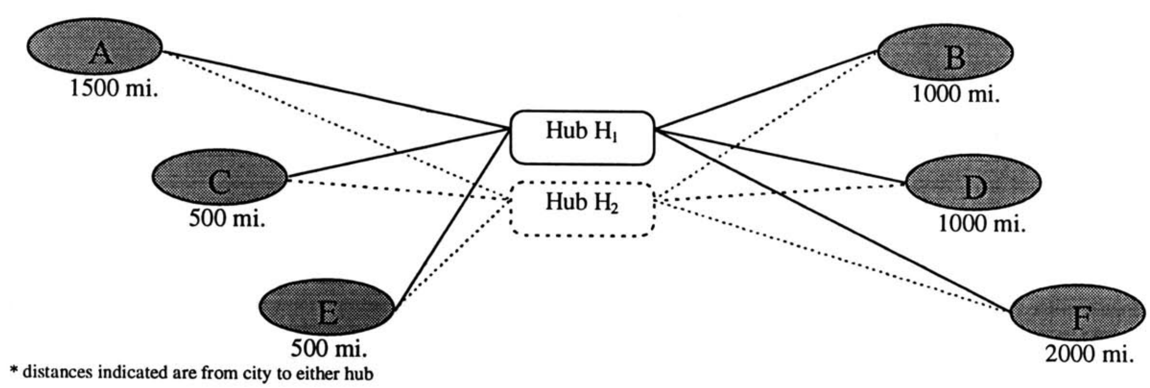

Figure 3.3: Spatial layout of PODS network

The spatial layout of the network itself (illustrated above in Figure 3.3) has also been designed to model a commonly occurring domestic US layout as best as possible. Of the six different leg distances from city to hub, two long-haul (one 2000-mile, one 1500-mile), two medium-haul

[61] Sell-up is the occurrence where a passenger for whom the desired fare class on a path is not available decides to purchase a higher fare class product for the same path.

(1000-mile), and two short-haul (500-mile) leg distances exist--this combination was chosen primarily because most major US domestic carriers offer a variety of short-haul and long-haul legs originating or terminating at their hubs. Furthermore, aircraft capacities on each leg have been chosen to be 100 seats, but such a determination is arbitrary since the demand itself in each O-D market can be scaled.

Finally, the demands on the legs were set in a systematic manner (although these demands are stochastic, an average and a standard deviation need to be set in the simulation). The methodology by which this was done has the short-haul legs and paths having the highest demands, while the longest-haul ones have the lowest demands (see Table 3.1 below).

Leg: Hub to/from city	**Leg Distance (mi.)**	**Leg Load (pax)**	**Local Pax (%)**
A	500	140	46%
B	500	125	40%
C	1000	110	41%
D	1000	90	39%
E	1500	75	33%
F	2000	60	33%

Table 3.1: PODS leg demands

The reasoning behind this approach is that higher demand on short-haul legs accentuates the need to perform O-D yield management, since the short-haul legs tend to have the lowest fares (especially in these scenarios, where a distance-based pricing scheme was used), and this is where demand "bottlenecks" tend to occur. Therefore, both facets of the O-D seat inventory control problem discussed in Section 1.1 become readily apparent. First, the tradeoff between accepting a lower fare class connecting passenger whose absolute revenue value is higher than that of a higher fare class local passenger is more uncertain. If the demand were highest on the longest-haul legs, the short-haul legs would no longer be creating these "bottlenecks" and we would usually accommodate most of the higher-fare connecting passengers. Second, the tradeoff between acceptance of two lower-fare customers whose combined revenue is higher than that of one long-haul, high-fare customer comes into play. If the highest demands were set to be in the longest O-

D markets, then the yield management system would again undoubtedly favor the long-haul customer and the revenue benefits of implementation would be less pronounced. It should also be noted that while the percentage of local passengers varies from 33% to 46%, the highest instances of local passengers are on the shortest-haul legs, for the same reasons as discussed above.

3.3 Other PODS input parameters

Aside from the various fixed market parameters described above, several other inputs are available within the PODS simulation and are described here (a complete listing is shown in Tables 3.2-3.4). Default values for system-level inputs which are constant in all simulations are listed next to the input parameters for informational purposes.

PODS version 7b System-Level Input Parameters	
• Number of airlines (2)	• Is passenger's first choice the only choice? (No)
• Number of markets (54)	• Elasticity multiplier (by pax type)
• Number of fare classes (4)	• Cancellation penalty ($0)
• Number of booking curves (2)	• No-show rate (0%)
• Number of observations used in forecaster (26)	• Leg fare calculation formula (for leg-based YM)
• Number of total samples (days) (600)	• Attributed cost k-factor (0.3)
• Number of legs (24)	• Time frame reoptimization days (see Figure 3.4)
• Number of passenger types (2)	• Primary, secondary z-factor (by pax type)
• Number of restriction categories (3)	• Cumulative booking probability by time frame
• Number of time frames (16)	• Cancellation rate (0%)
• Number of samples burned (200)	• Last time frame where fare classes are available
• Number of trials (20)	• Passenger type k-factor (0.4)
• System k-factor (0.1)	

Table 3.2: PODS version 7b system-level input parameters

PODS version 7b Airline Input Parameters	
• Seat optimization (YM) method (see Section 2.3)	• For HBP/Netbid, number of bookings between availability processor executions
• Forecasting method (see Section 2.1)	• For HBP/Netbid, is the forecast believed?
• Detruncation method (see Section 2.2)	• For DAVN, displacement method
• For projection detruncation, # iterations	• For DAVN, virtual class revision scheme
• For projection detruncation, minimum number of unclosed observations	• Number of virtual classes + boundaries
• For projection detruncation, quitting criteria	• Theft/standard nesting used
• For projection detruncation, tau value	• Bid price scaling, calculation constant
• Sigma scaling for booking capacity	• Probability of sell-up

Table 3.3: PODS version 7b airline input parameters

PODS version 7b Market-Level Input Parameters	
• Market k-factor	• Base fare in market
• Capacity on each leg	• Market demand
• Distance of each leg	• Percent of business passengers in market (by pax type)
• Mean schedule tolerance	• Preferred airline probability in market
• Time of day curve	• Fares by market/fare class (i.e., ODF)
• Number of paths in market	• Path departure/arrival time
• Distance of market	• Path quality index
• Market delta-T	• Number of legs traversed by path
• Denied boarding penalty	

Table 3.4: PODS version 7b market-level input parameters

First, the number of observations used from the historical database in projecting demand estimates is 26 (i.e., the 26 most recent flight observations). This number is sufficiently large to ensure that single observations that are unusually high or low will not have a pronounced effect on the forecasts; and as PODS has been developed to model a stationary process in which seasonality and other airline demand trends are disregarded, such effects need not be considered or adjusted for. Second, for a given flight, there are 16 time frames (booking intervals), occurring at semi-regularly spaced periods; at the beginning of each time frame is when the seat optimization/allocation routine is performed. In our experiments, from 63 days before departure until 35 days before departure, booking limits are updated weekly, so the length of the first 5 time frames is 7 days. Then, from 35 days before departure until 7 days before departure, booking limits are updated bi-weekly, hence, the length of the next 8 time frames is 3.5 days (in the graphical representation below integers are used for ease of illustration, so one 3-day interval and one 4-day interval occur in each week). Finally, in the last week before departure (i.e., 7 days before departure until 1 day before departure), booking limits are recalculated every 2 days, so the length of the last 3 time frames is 2 days. A graphical depiction of the time frames is shown in Figure 3.4.

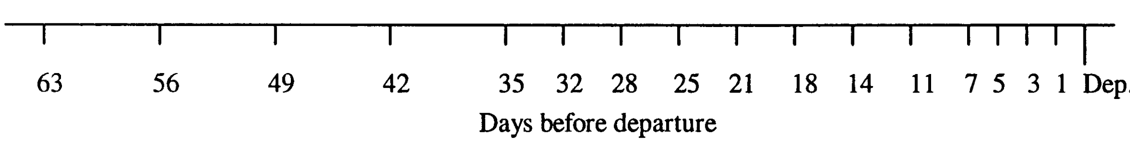

Figure 3.4: Time frame intervals for a flight departure in PODS

Third, the fare classes themselves are also able to be specified, but for consistency in our set of simulation cases, 4 distinct fare classes will be used (labeled Y, B, M, and Q). The fare magnitudes offered in these fare classes vary by O-D market and although in normal airline practice fares are assigned completely on an O-D market basis, in PODS they have been set to have some correlation to the distance traveled. A base set of fares offered is \$800, \$400, \$300, \$200 for Y-, B-, M-, and Q-class, respectively; this being for the 2000-mile O-D market. Fares for other O-D markets are then assigned with the Q-class "base" fare being multiplied by a ratio of 1.6 for each doubling of distance (or multiplied by 1/1.6 for each halving of distance), while the relative ratios of the Y-, B-, and M-class fares to the Q-class fare remain constant at 4, 2, and 1.5, respectively. The complete fare schedule for the network is illustrated in Table 3.5 below.

	O-D Market Distance (mi.)							
	500	**1000**	**1500**	**2000**	**2500**	**3000**	**3500**	**4000**
Y	\$320	\$500	\$660	\$800	\$920	\$1020	\$1205	\$1280
B	\$160	\$250	\$330	\$400	\$460	\$510	\$602.5	\$640
M	\$120	\$187.5	\$247.5	\$300	\$345	\$382.5	\$451.88	\$480
Q	\$80	\$125	\$165	\$200	\$230	\$255	\$301.25	\$320

Table 3.5: PODS version 7b fare schedule

The individual fare products within any given O-D market are differentiated in that there are up to three different restriction categories imposed on each fare product, as passengers inherently assign some corresponding dollar value to each restriction. The fare classes and their associated restriction categories are detailed in Table 3.6 below (restrictions on a fare product are shown as shaded blocks).

Fare class	Advance purchase requirement	Stay requirement	Cancellation/ change penalty	Refundability
Y	--none--	--none--	--none--	yes
B	7 days	Sat. night	--none--	yes
M	14 days	Sat. night	yes	yes
Q	21 days	Sat. night	yes	no

Table 3.6: Ticket restriction categories

With some of the seat optimization routines (namely, VEMSRb, DAVN, and HBP), a virtual fare class bucket arrangement is necessary (virtual buckets are described in Section 2.3.1.2). Within PODS, such an arrangement has been set to be comprised of 6 virtual fare classes, where each ODF on each leg maps into one of the associated virtual fare classes. The fare divisions among the six virtual buckets are shown below in Table 3.7 (these are for the VEMSRb and HBP seat optimization algorithms; DAVN makes virtual bucket divisions based on the calculated pseudo fares--see Section 2.3.1.4).

Virtual fare class	Fare value
Y1	< $1300
Y2	< $700
Y3	< $420
Y4	< $310
Y5	< $215
Y6	< $140

Table 3.7: Virtual fare class boundaries for VEMSRb, HBP

With the virtual classes, the fare disparities among the different fare classes in the individual O-D markets are accounted for by the groupings which occur. For example, combining Table 3.5 and Table 3.7 it can be seen that both the Y-class fare for the 500-mile market and the Q-class fare for the 4000-mile market map into the same virtual bucket (Y3)! Methods using virtual classes are therefore attempting to distinguish seat availability by *fare class value* rather than by *fare class designation*.

There are also a variety of other input parameter effects which are able to be tested in the PODS routine, but will not be explored in this thesis. Some of these include no-show percentages (i.e., the percent of passengers who reserve but do not show up for the flight), overbookings (i.e., the margin by which excess bookings are taken on a flight in order to account for expected no-shows), and denied boardings (i.e., when flights are overbooked but no-shows are lower than expected, some fraction of passengers are denied boarding, either voluntarily or involuntarily, on the given flight). Others include inputs of the demand variability in the system or in a market; for example k-factors, by which the standard deviation of demand for a flight sample is proportional to its mean, or z-factors, for which the variance of the demand random variable is proportional to its mean.

In addition to the many fixed parameter inputs which have been described so far, there are also a number of input parameters which will be changed within different simulations. The first three were discussed previously in Chapter 2. First, there are five choices for the actual seat optimizer itself, described briefly in Section 2.3.1. Second, the forecasting routine will be varied; these variations are described in Section 2.1. Third, variations in the detruncation method will also be tested; these are described in Section 2.2. Other input parameters to be varied within the scope of this thesis are discussed in the subsections below.

3.3.1 Demand factor changes

One of the input parameters which will be varied within this thesis is that of demand factor. Although the demand generated in PODS is stochastic, the demand factor is defined as the ratio of the *average* realized demand to the aircraft capacity (which was set to equal 100). Hence, a demand factor of 1.0 means that demand, on average, will be equal to 100 passengers, while a demand factor of 1.2 means that demand will be 120 passengers on average. The testing of demand variations is of interest to see how the system revenues are impacted not only by forecasting, detruncation, and seat optimization changes, but also whether these results are uniform under higher or lower demand scenarios.

We would expect that under low demand scenarios there should not be as much variation of the network revenue or loads, simply because seat optimization limits are rarely being pushed to capacity by demand and therefore fewer passengers are being spilled, so the resultant effect of implementation of a less effective forecasting routine or seat optimization algorithm will be less pronounced. However, under high demand scenarios these effects should be accentuated, since the opposite case is occurring; fare class capacity is often being reached, so booking limit "decisions" are constantly being made by the seat optimizer as to whether or not to spill high-fare or low-fare customers. Because of this, more accurate or effective forecasting measures will be essential to better revenue performance. In simulation experiments run in this thesis, three different demand factor parameter values will be used: 0.8, 1.0, and 1.2. A demand factor of 0.8 simulates a low-demand scenario as explained above, while a demand factor of 1.0 simulates an average demand scenario and a demand factor of 1.2 simulates a high-demand scenario. These demand factors chosen correspond roughly to average system load factors of 70%, 78%, and 83%, respectively, which are rather plausible results based on general airline industry estimates.

3.3.2 Competitor airline parameter changes

In contrast to the studies referenced earlier, the version of PODS being used in this thesis (version 7b) has the capability of permitting not only parameter changes for the airline of interest, but also for the competing airline. That is, effects of different combinations of forecasting method, detruncation method, and seat optimization routine for the competing airline can be tested. The first group of results examined (Section 4.1) will analyze forecasting, detruncation, and seat optimization changes by the airline of interest versus a common set of parameters where the competing airline is using EMSRb seat optimization (see Section 2.3.1.1) with pickup forecasting and booking curve detruncation. This is a realistic scenario, as the EMSRb seat optimizer is a commonly-used technology at many airlines both in the domestic US and internationally (pickup forecasting and booking curve detruncation can also be thought of as the "base case" choices for these methods). Therefore, the resultant impacts to the airline being studied (i.e., not the competition) against these scenarios are of importance, as they can be readily applied to realistic situations in the industry.

However, also of interest are scenarios in which the competition has a different, more advanced yield management algorithm, so insight can be gained into what impacts come about from forecasting or seat optimization changes by the airline being studied. That is, we are interested not only in competitive situations where the airline being studied has taken the yield management initiative, but also in a myriad of possibilities where the competition has also implemented a more advanced set of yield management tools (a case which occurs among the larger airlines in the US industry). This allows us to move toward achieving the second goal of the thesis; namely, determining which combinations for the airline of interest are most beneficial for a wide variety of scenarios under which the competition operates.

3.4 Network vs. single-leg optimization

So far, related previous PODS studies have concentrated on the revenue effects of two airlines competing on a single leg only. However, this thesis will extend this notion to the network case, where two airlines compete over a network of cities connected by a hub airport. In the single-market case of two competitors serving one origin and one destination, seat optimization routines are looking simply at the marginal benefits and costs of additional passengers in a given fare class on the route; even so, Wilson[62] has shown that implementation of an EMSRb seat optimizer can result in revenue gains on the order of 8-10% over the option of using first-come, first-served booking (i.e., no capacity restrictions on the fare classes), while the corresponding loss to the carrier who retains the first-come, first-served approach is on the order of 2%. Other scenarios such as those with a dominant carrier and a weaker carrier--where one carrier offered more daily flights than the other--were tested, for which the revenue difference was much more disparate (on the order of 25%). However, in all simulations performed here, equal prices, frequencies, and image factors for both carriers will be assumed.

When the yield management concept is extended to the multiple-origin, multiple-destination network case with not only leg-based seat optimization routines but also more advanced O-D seat optimizer choices available (i.e., those which consider factors such as downline displacement costs of passengers and the effects of local vs. connecting passengers), the implementation of a yield management system should also be expected to provide moderate revenue gains. However, a good question arises as to how much of the enhanced revenue comes about due to actual implementation of the seat optimizer, and how much can be attributed to the forecasting/detruncation method used in the simulation. Skwarek estimated that in the symmetric single-leg cases examined by Wilson[63], the proportion of revenue increase attributable to the seat optimizer itself is on the order of 65% (the other 35% resulting from a better detruncation method)[64]. In the multiple-leg case, it is also of interest to see whether revenue increases due to

[62] Wilson (1995). Section 5.1 analyzes symmetric two-path scenarios where one or both carriers can implement the EMSRb seat optimizer or use the traditional first-come, first-serve option.
[63] Wilson (1995).
[64] Skwarek (1997), Section 7.2.2.

the implementation of improved O-D seat optimizers of about 0.5-2% found by Lee[65] (note that this is over a base case of EMSRb as opposed to first-come, first-served as was done in Wilson's single-leg case) are attributable to forecasting or detruncation improvements in the same ratio as before.

3.5 PODS output to be examined

When the yield management simulations are being run, some method of comparison needs to be defined to allow for consistent measurement of the relative performance of different combinations of forecaster, seat optimizer, etc. There are primarily four measures by which such performance will be measured: (1) overall network revenue; (2) passenger leg loads; (3) fare class closures; and (4) forecasted remaining demands.

First, overall network (system) revenue is just the average network revenue over the 20 trials, with the network revenue for each trial being the sum over all O-D markets of the products of each O-D passenger in that market and their corresponding fare. It is a good measure of the general performance for the airline as a whole, especially because in reality, forecasting or seat optimization systems such as those being tested are implemented by an airline on a system-wide basis (and not usually in just a certain market or on a certain flight leg). Additionally, total system revenue provides an average, unbiased measure of the network performance of the quantity we are looking to maximize (i.e., revenues).

Second, passenger loads grouped by flight leg is a measure which provides more insight into why revenues may have increased or decreased, as trends such as more high-fare or low-fare passengers being taken on or spilled can be seen directly. For example, system revenues may increase, but for two completely different reasons--one possibility being a better seat protection algorithm in which total loads are lower but more high-yield passengers are taken (whose incremental revenue contribution more than offsets the revenue lost by the low-yield passengers no longer taken), another being that overall loads are higher as individual loads in each of the

[65] Lee (1998). In Section 4.1.3, values in this range were obtained for the different seat optimization methods at a demand factor of 1.0 under three different network scenarios (i.e., 4 spoke cities with 1 common hub, 4 spoke cities with 2 decoupled hubs, and 6 spoke cities with 2 decoupled hubs).

different fare classes are higher. Therefore, examination of the passenger leg loads provides a better understanding of the direct impact on the revenue increase or decrease seen by the first performance measure discussed.

Third, fare class closures will be used to analyze the relative availability of the different fare classes on average. For any given ODF, the time frame in which the fare class closed can be determined (this is the point during the booking process after which any additional passenger requests are spilled); the more available an ODF, the higher passenger loads we would expect, which can be either good or bad depending on whether it is for a high or low fare class and how large the potential demand is for the other fare classes is.

Finally, forecasted remaining demands are insightful in that they are the driving force behind what influences the fare class closures, passenger loads, and system revenues. Higher forecasts for an ODF will generally lead to higher protection limits, more fare class availability, and therefore higher loads for that particular ODF. As indicated, such performance measures are useful in attempting to explain the different phenomena which can result from variations in any of the PODS input parameters--these measures will next be discussed in the context of actual simulation runs in Chapter 4.

Chapter 4

Analysis of Forecasting/Detruncation Method Changes

While the first three chapters have provided a background on forecasting and detruncation algorithms (as well as yield management in general) along with descriptions of the different methods for such algorithms employed by the PODS simulation and the functionality of the PODS system itself, this chapter concentrates on the testing and analysis of different cases of forecasting, detruncation, and seat optimization within the yield management framework of PODS. This is done to provide insight into the relative advantages and disadvantages of each particular method or combination of methods. In all tests performed, PODS version 7b was used; details of which were described in Chapter 3.

First, base case yield management routines will be analyzed, in which Airline A changes only the seat optimizer while base case yield management is used by Airline B (i.e., the forecaster, detruncator, and seat optimizer are held constant at the base case)[66]. Comparisons will then be made between these base cases and those in which Airline A has the ability to change any of these parameters (i.e., the forecaster, detruncator, and seat optimizer will be varied by Airline A yet will still be held constant at the base case by Airline B). Second, the impacts of "competitive scenarios" will be studied--these are situations in which both airlines match each other's seat optimization algorithm but employ different forecasting/detruncation method combinations. Third, alternative yield management base case combinations (i.e., those where seat optimizers other than EMSRb are used by Airline B) with variations in forecasting and detruncation method by either one or both airlines will be investigated, in order to determine the relative competitiveness of different combinations of methods within the yield management context. The

[66] The airline whose results are of interest will heretofore be referred to as "Airline A", while "Airline B" will be used in reference to the competing airline.

subsequent section analyzes changes in the parameter τ, used in the projection detruncation algorithm, to determine whether such changes alter the results observed in the preceding sections. The final section looks to explain the factors contributing to the resultant revenue increases or decreases seen in the preceding sections, along with a synopsis of which forecasting/detruncation methods perform best given a particular seat optimization algorithm. Comparisons will also be made with earlier studies, and reasons behind any similarities or differences will be explained.

Although the primary criterion for measuring the relative performance of the forecasting and detruncation methods in this thesis will be overall system revenues[67] as it provides a good macro-level measure of the revenue values obtained by the system, others will also be used. Namely, passenger loads (both actual and forecasted) and fare class closures will be examined to determine the relative performance of the methods and the reasons for any differences; additionally, comparisons will be made with similar studies under earlier PODS versions where appropriate. While the PODS output gives absolute network revenue values, the relative percent increase in revenue as compared to some base case set of methodologies is of interest, and will always be computed relative to a base case of pickup forecasting and booking curve detruncation, with a given seat optimization method. Lastly, a final note is made about the nomenclature which follows in the subsequent sections. For simplicity on the graphs, the forecasting and detruncation combinations are abbreviated as follows: PU = pickup forecasting, R = regression forecasting, BC = booking curve detruncation, and P = projection detruncation.

4.1 Impacts against base case yield management

A good starting point for analysis of different forecasting and detruncation schemes in the PODS simulation is that of scenarios which compare results when changes are made by one airline competitor while the other uses a static "base case" combination of forecaster, detruncator, and seat optimizer. This allows a controlled experiment in which the direct impacts of changes in these methods can be isolated and analyzed by varying the forecasting/detruncation method. Revenue results can be examined, along with passenger loads, forecasts, and even fare class closure times. In the ensuing discussion, results will be primarily of interest for the airline

[67] System revenues as calculated are the airline's network revenues averaged over the 20 trials.

implementing forecasting, detruncation, and seat optimization changes, designated as Airline A. Airline A will therefore have a choice among the various combinations of forecasting, detruncation, and seat optimization methodologies as presented in Chapter 2. Airline B will be designated as the competing airline; it will initially use a base case combination of EMSRb seat optimization methodology with pickup forecasting and booking curve detruncation[68]. Furthermore, relative comparisons will be made against either the full base case scenario (i.e., both airlines using EMSRb with pickup forecasting and booking curve detruncation), or a modified base case, in which Airline A is using pickup forecasting and booking curve detruncation along with a different seat optimization routine (Airline B still using EMSRb with pickup forecasting and booking curve detruncation).

What do we expect from the base case results? Among the seat optimization methods, we would expect that EMSRb performs the worst, as its seat protection calculations are based solely on fare classes by flight leg, with no distinction between fare classes in high- and low-fare markets, and no consideration for passenger displacement costs. VEMSRb attempts to correct for the former problem by grouping the ODFs into virtual fare buckets, but it still does not account for passenger displacement costs and should therefore perform better than EMSRb, but worse than Netbid, DAVN, or HBP. These last three methods all attempt to adjust for these problems and should theoretically perform better than either EMSRb or VEMSRb in all cases. As to the relative rankings among these three methods, that is yet to be determined (Section 4.1.1), as are the effects of changing the forecasting and detruncation routine in conjunction with the seat optimization algorithms (Section 4.1.2).

4.1.1 Base case forecasting and detruncation

The base case forecasting and detruncation results are important in that until now, all yield management simulations of a multiple-leg network run under PODS have assumed a base case of pickup forecasting and booking curve detruncation. Therefore, such results detail what has been previously determined about the relative performance of the various seat optimization routines

[68] EMSRb is the most basic of the seat optimization routines tested, and it is currently used by many airlines worldwide in revenue management applications. Most current airline applications of this EMSRb methodology (as well as other seat optimizers) use simple pickup forecasting and booking curve detruncation.

tested in earlier versions of PODS. Hence, a good starting point is a reproduction of similar results obtained by Lee[69], which illustrate the relative performance of the various seat optimization routines tested (forecasting and detruncation were always held at the base case of pickup and booking curve, respectively). Table 4.1 below shows the network revenues by seat optimization routine for Airline A under each of three different demand factors, along with the absolute increase over the base case revenues under EMSRb. It can be seen that for all demand factors tested, all seat optimizer changes from the base case of EMSRb produce revenue increases, and DAVN performs the best by providing the largest revenue gain.

	Airline A System Revenue					
Seat Optimizer	***DF 0.8***	***Incr. over EMSRb***	***DF 1.0***	***Incr. over EMSRb***	***DF 1.2***	***Incr. over EMSRb***
EMSRb	189,813	---	226,954	---	259,762	---
VEMSRb	191,115	+1302	230,529	+3575	263,265	+3503
Netbid	190,076	+263	228,871	+1917	262,526	+2764
DAVN	191,365	+1552	232,231	+5277	266,801	+7039
HBP	191,016	+1203	231,414	+4460	266,292	+6530

Table 4.1: Total system revenues for Airline A

To gain more intuitive insight into the relative performance of the seat optimization routines and to furthermore be able to draw comparisons across demand factors, relative percentages can be computed between the revenue values listed in Table 4.1. In Figure 4.1, a percent revenue increase is shown, where this percent increase is compared to the base case of Airline A using EMSRb seat optimization with the default forecasting and detruncation routines (i.e., pickup and booking curve, respectively). Note that Airline B is still using this same base case in all scenarios; that is, only the seat optimization routine of Airline A is being varied under each demand factor.

[69] Lee (1998). Refer to Section 4.1.3, although it should be noted that PODS version 6 was used for Lee's results.

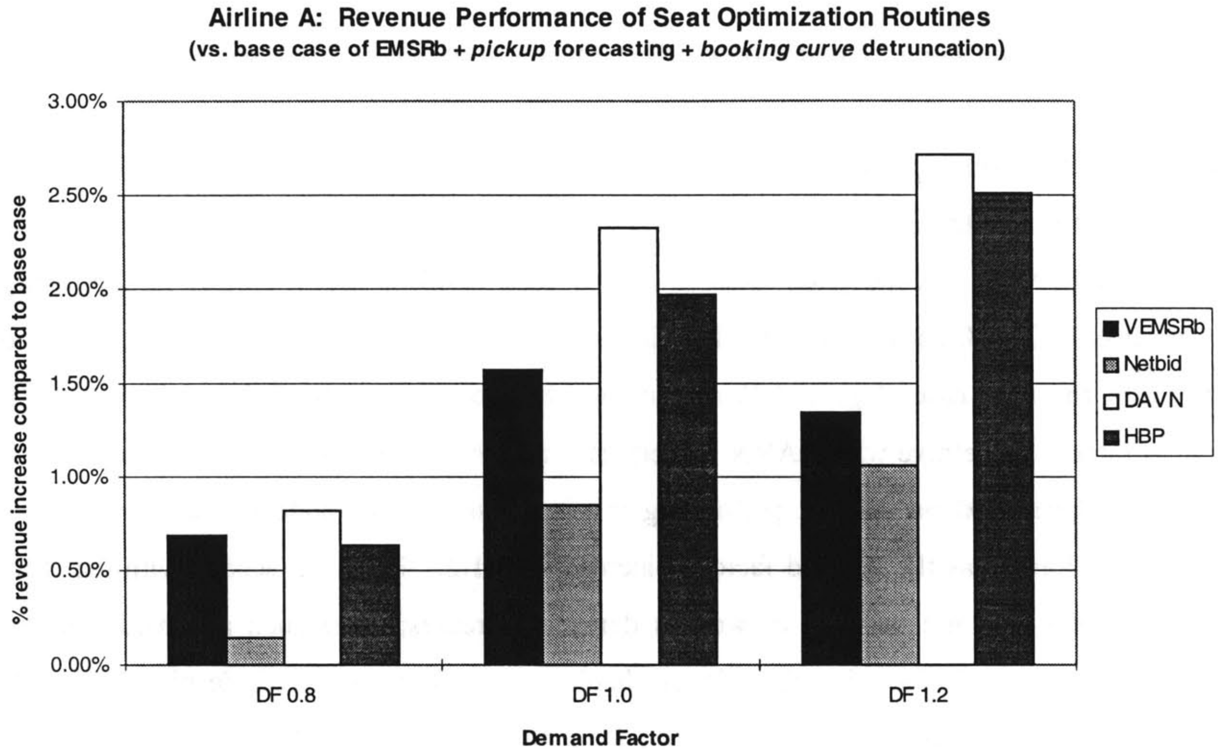

Figure 4.1: Revenue performance of seat optimization routines for Airline A

Basically, from Figure 4.1 above, we can make inferences about the performance of these alternate seat optimization algorithms relative to the EMSRb base case and to one another. It can be seen that under all demand factors, DAVN results in the highest percentage gain compared to the base case, while Netbid produces the lowest gain (although it still has higher revenues than EMSRb). Additionally, as the demand is higher (demand factor increases), the relative gains from improvement of the seat optimizer also increase, from an average of 0.57% in the DF 0.8 case to averages of 1.68% and 1.91% in the DF 1.0 and DF 1.2 cases, respectively. This is to be expected, since higher demand only accentuates the benefit of O-D yield management, as more passengers (and consequently, higher-fare passengers) are willing to travel and there is greater revenue benefit from distinguishing between O-D paths and fare classes. When the demand is low as in the DF 0.8 case, the seat protection limits are rarely being reached (or the bid prices are too low), and little traffic is therefore being turned away; in essence, the yield management is not functioning to its potential. However, it is interesting to note that even in the low demand case

(DF 0.8), there are still some benefits resulting from implementation of a "better" seat optimization routine when the competitor airline remains with the base case scenario.

In terms of the relative performance of the seat optimization routines themselves, several items are noteworthy. In the case of low demand (DF 0.8), VEMSRb, DAVN, and HBP all perform about equally, with revenue gains of 0.63-0.82% over EMSRb, while Netbid is substantially lower, at only 0.14% above EMSRb. For DFs 1.0 and 1.2, we see that DAVN performs best relative to the base case (Figure 4.1) and in overall system revenues (Table 4.1), while HBP remains rather competitive with DAVN, underperforming it by only about 0.2-0.3% in each case. VEMSRb places next on the list, performing relatively well in the moderate demand case, but being outperformed as the demand factor is increased to 1.2. This is in accord with the theory that VEMSRb does not perform as well as demand increases; since such a "greedy" scheme favors the longer-haul itineraries, these longer-haul passengers are accepted, despite the fact that higher revenue could result from carrying two local passengers (given that the demand is rather high) instead of a single connection passenger. Finally, Netbid provides the lowest relative revenue percentage gain; in fact, DAVN's percentage gains outdo those of Netbid by a factor of approximately 2.5! This is a somewhat surprising result, as one would expect that Netbid should perform almost as well as DAVN or HBP, since it is accounting for passenger displacement costs as the other methods do. Investigations into why this occurs will be one of the major topics of the next section. More detailed analysis of these results are provided by Lee[70].

Several basic conclusions which form the basis for later comparisons in this thesis can be drawn from these preliminary results. First, the percentage revenue increase for Airline A of moving from EMSRb to one of the other seat optimization routines (even using base case forecasting and detruncation) when the competitor uses a complete base case (i.e., EMSRb, pickup, and booking curve) are on the order of 1.5% in the medium demand case. Second, under the base case forecasting and detruncation methods for the demand factors analyzed, Netbid performs rather poorly, while DAVN (and HBP) perform very well in terms of revenue. The poor performance of

---

[70] Lee (1998), Section 4.1.3. Netbid was initially found to have very poor performance on the very small four-city (with one common hub) network, as the gap between bid prices was too large to have a good passenger mix. The six-city results provided a small improvement, although Netbid's relative performance was still quite poor.

Netbid is in contrast to what was expected; intuition would lead us to believe that Netbid outperform VEMSRb and be on par with the other methods that account for passenger displacement costs (DAVN, HBP). Finally, the trends among the rest of the methods (with the exception of Netbid) do follow what was expected, as discussed above in Section 4.1.

4.1.2 Forecasting and detruncation changes by a single competitor

As seen in Section 4.1.1 above, our initial hypothesis about the relative performance of the seat optimization methods was mostly correct, with the exception of explaining the poor performance of Netbid. A good question therefore arises in the analysis of why Netbid's performance is substantially lower than the other methods with which it should be competitive. Several conjectures have been put forth in this regard in attempting to explain these results--the fact that Netbid is solving a deterministic linear program rather than a probabilistic one, the small size of the network on which it is being tested, and the incompatibility of the forecasting and/or detruncation method. The latter possibility--that the base case forecasting and detruncation routines used are not as effective in providing the required demand inputs to Netbid's seat optimization routine--will be examined in depth in this section. This will be done by examining the general trends occurring when forecasting and detruncation methods are changed by Airline A, while the competing airline still uses the same base case as in the preceding section (doing so allows the effects of the forecaster and detruncator to be isolated).

4.1.2.1 System revenues

As in Section 4.1.1, system revenues will first be compared to analyze the general trends among the combinations of seat optimizer, forecaster, and detruncator tested. Figure 4.2 presents the percent increase in system revenues resulting from changes in the seat optimization routine and forecasting/detruncation method combination for Airline A at a demand factor of 1.0, where the percentage gains are compared to the base case system revenues for Airline A using EMSRb with pickup forecasting and booking curve detruncation.

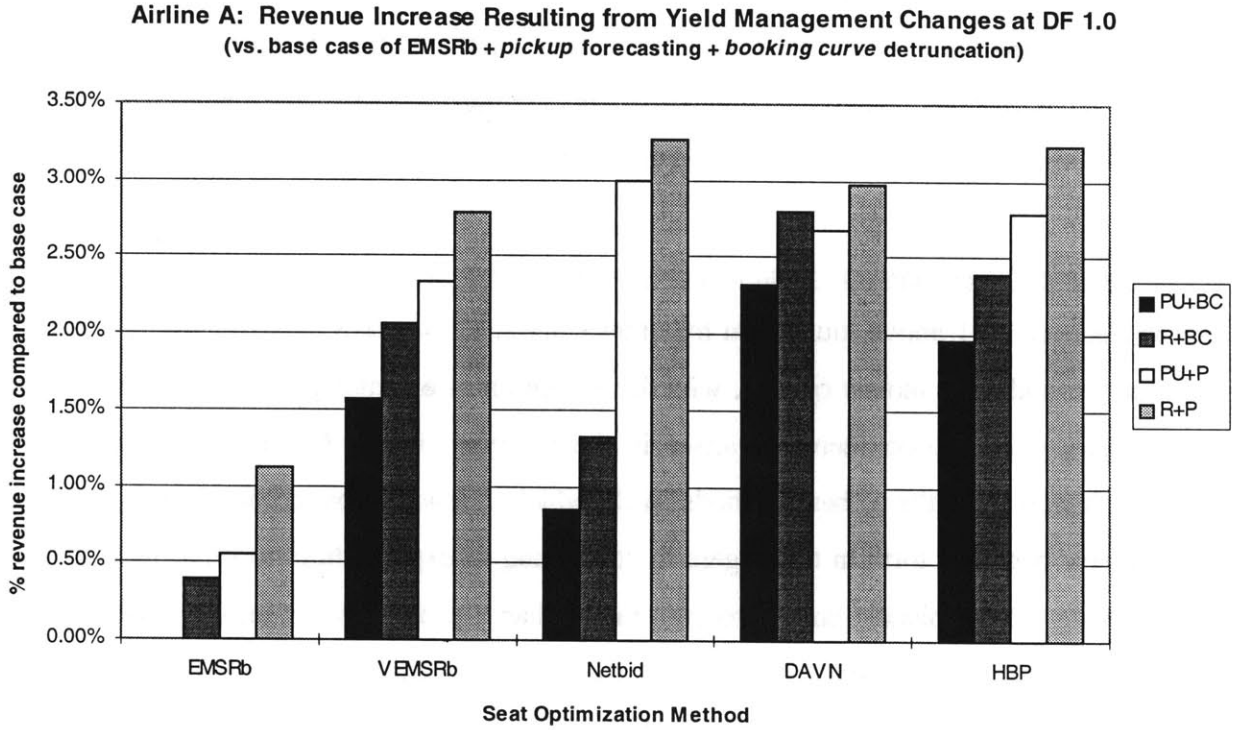

Figure 4.2: Revenue increases from yield management changes at DF 1.0

Two clear inferences can be made from this graph. First, the relative trends of forecasting and detruncation combination are evident, regardless of the choice of seat optimization routine. That is, the "base case" combination of pickup forecasting and booking curve detruncation performs worst, while regression forecasting and booking detruncation perform slightly better, pickup forecasting and projection detruncation are even better, and regression forecasting and projection detruncation provide the largest increase over the base case[71]. The relative magnitudes of these percentage revenue gains depend on the choice of seat optimization. DAVN provides gains of roughly 2.5% when booking curve detruncation is used, HBP follows second with gains of approximately 2.25%. Of interest are the relatively small increases for Netbid of around 1.0%--only outperforming the very small gains achieved with EMSRb. However, under projection detruncation the percent increase in revenues are much higher, with Netbid performing best at more than a 3.0% gain. DAVN and HBP have comparable performance in this scenario, at about 2.75% and 3.0%, respectively.

[71] Only for DAVN with pickup forecasting and projection detruncation is this trend violated.

Second, the volatility of the seat optimization routines with respect to the forecasting and detruncation method combination is evident. The EMSRb, VEMSRb, and HBP seat optimization routines are moderately affected by forecasting and detruncation changes (on the order of a 1.25% difference between the maximum and minimum values--see Table 4.2); DAVN is minimally affected (0.64% difference), and Netbid is highly volatile (2.43% difference). Revenue performance also followed similar relative trends in the DF 0.8 and the DF 1.2 cases.

Netbid's revenue performance was far below what was expected when the base case forecaster and detruncator were used in Section 4.1.1; however, it exceeds all other seat optimization algorithm choices when the combination of pickup forecasting and projection detruncation or regression forecasting and projection detruncation is implemented. It therefore appears that the driving force behind Netbid's revenue disparity is the choice of detruncation method. Simple booking curve detruncation, while performing adequately for other seat optimizer choices, does not perform well for Netbid; however, when projection detruncation is chosen, Netbid's performance falls into line with the other similar seat optimization methods. Similar trends were also seen in the DF 0.8 and DF 1.2 cases, illustrated in Figures 4.3 and 4.4, although one can see that the percent revenue increases definitely become larger with increases in demand factor.

Table 4.2 illustrates these trends for the different demand factors tested. The first category of data is the largest percentage difference between the various forecasting and detruncation method combinations, under each of the different seat optimization algorithms. Such a comparison provides information about the volatility of the seat optimizer with respect to the combination of forecasting and detruncation changes--the larger this percentage, the more sensitive the seat optimizer is to the forecasting and detruncation methods. The shaded blocks are the values for Netbid under the different demand factors; as it can be seen, Netbid's variance is quite high relative to the other seat optimization methods (more than double any other seat optimization method, with the exception of VEMSRb in the high demand case). It should also be noted that DAVN is the most robust method, as the variance is only on the order of 0.6%, regardless of the demand factor.

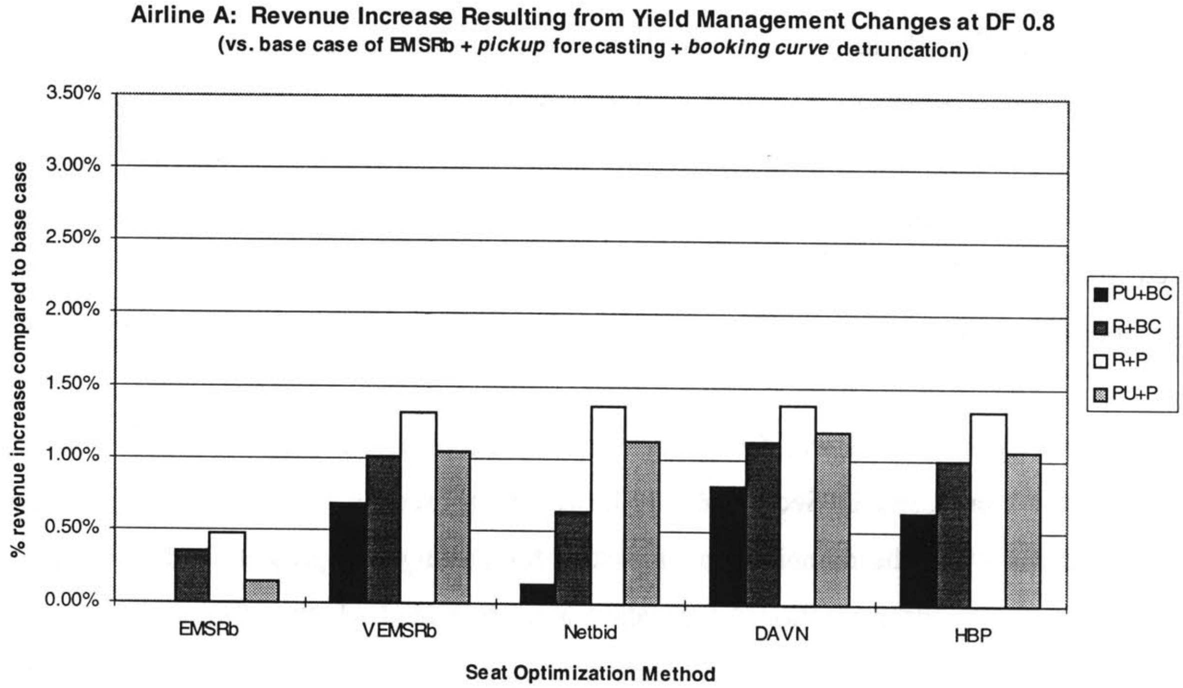

Figure 4.3: Revenue increases from yield management changes at DF 0.8

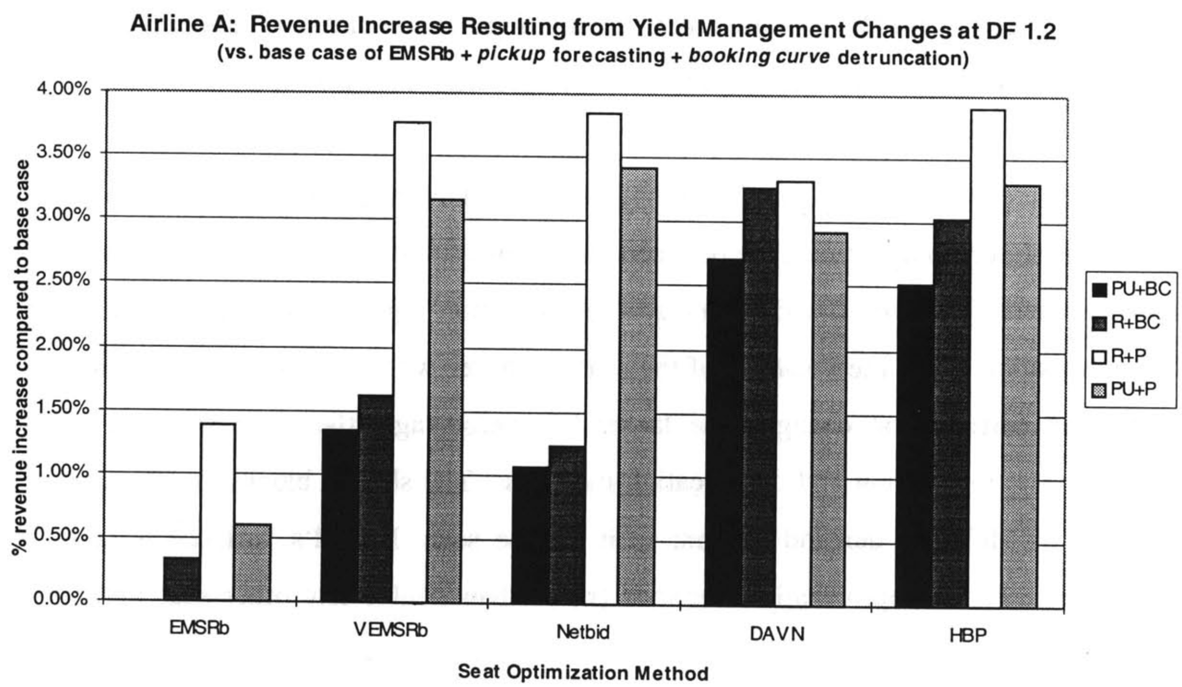

Figure 4.4: Revenue increases from yield management changes at DF 1.2

	DF	EMSRb	VEMSRb	Netbid	DAVN	HBP
				Seat Optimization		
All Combinations Max - Min	**0.8**	0.47%	0.62%	**1.23%**	0.56%	0.70%
	1.0	1.12%	1.21%	**2.43%**	0.64%	1.27%
	1.2	1.38%	2.41%	**2.78%**	0.62%	1.38%
Avg. Increase from Forecasting Change	**0.8**	0.34%	0.30%	0.37%	0.25%	0.32%
	1.0	0.48%	0.47%	0.38%	0.38%	0.44%
	1.2	0.56%	0.44%	0.30%	0.48%	0.56%
Avg. Increase from Detruncation Change	**0.8**	0.13%	0.33%	**0.86%**	**0.31%**	0.38%
	1.0	0.64%	0.74%	**2.05%**	**0.26%**	0.83%
	1.2	0.82%	1.97%	**2.49%**	**0.14%**	0.83%

Table 4.2: Relative percentage changes among forecasting/detruncation combinations

The second category of data in Table 4.2 is the average percent change in revenues resulting from a change in the forecasting method from pickup to regression (averaged over the two detruncation methods). This comparison details the sensitivity of the seat optimization routines to the forecasting method alone. Here, we see that only slight increases result in all cases (the maximum increase is only 0.56%), indicating that large impacts are not felt by the seat optimization routines when their forecasting method is changed. The data above also allows us to infer that although the two forecasting routines do not produce widely disparate results, regression forecasting does indeed provide slightly higher revenues relative to pickup forecasting in all cases.

The third and final category of data in Table 4.2 is the average percent change in revenues resulting from a change in the detruncation method from booking curve to projection (averaged over the two forecasting methods); this provides insight into the volatility of the seat optimization routines to the detruncation method alone. Unlike the small changes in forecasting method, we see here that changing the detruncation method to projection from booking curve has varying impacts, depending on the seat optimizer choice. The shaded blocks depict the percentage gains due to a change from booking curve to projection detruncation for Netbid (lightly shaded) and DAVN (darkly shaded); for Netbid they are quite high (a change of 2.0-2.5% for DFs 1.0 and 1.2), while for DAVN they are rather low (only about a 0.3% change or less). The other seat optimization algorithms experience moderate increases of approximately the same magnitude as the increases due to forecasting changes; all are less than 1.0% with the exception of VEMSRb in the DF 1.2 case.

Another set of comparisons can be made using the same revenue data as before, where the data is transformed in a different way. This is done by still examining percentage revenue increases, but this time comparing them to a base case of the same seat optimizer using pickup forecasting and booking curve detruncation (e.g., Netbid with regression and projection would be compared to a base case of Netbid with pickup and booking curve). What this allows is an interpretation of the revenue gains due *solely* to the change in the forecasting and detruncation method under each seat optimization algorithm, as opposed to the gains resulting from the particular combination of methods.

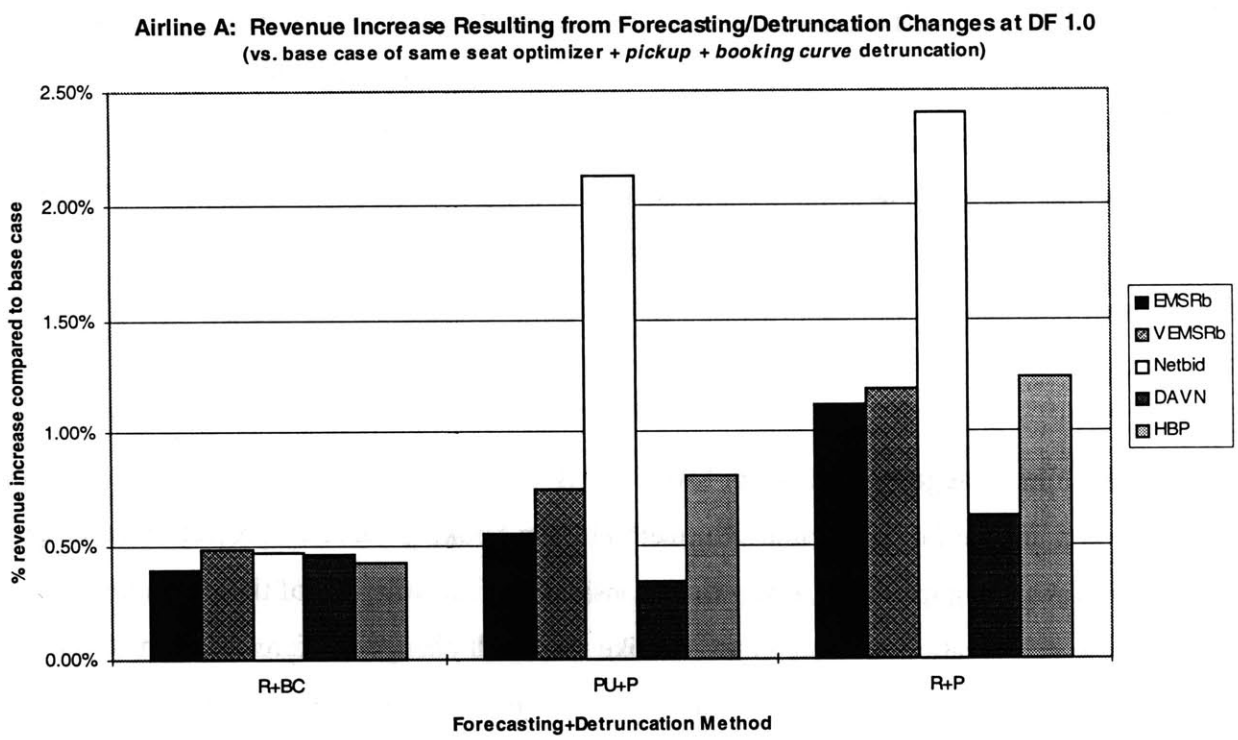

Figure 4.5: Revenue increases from forecasting/detruncation changes

The average percentage gains using a base case of the same seat optimizer with pickup forecasting and projection detruncation are presented in Figure 4.5 above. For a change to regression and booking curve, a 0.45% average revenue increase is experienced; a change to pickup and projection causes an average increase of 0.92%; and a change to regression and projection results in a 1.32% average increase. These values confirm the same relative performance as compared with the full base case of EMSRb, pickup forecasting, and booking curve detruncation examined earlier. Also corroborating the same results as the full base case are the relative performances of Netbid and DAVN. It can be seen that Netbid under projection detruncation has revenue percentage increases of 2.13% and 2.41% for pickup and regression forecasting, respectively--much higher than any of the other methods. On the other hand, DAVN's percentage increases are somewhat lower than average; this is due to the good performance of DAVN under the base case (i.e., DAVN with pickup forecasting and booking curve detruncation performed quite well; hence, the changes in detruncation method will have much less pronounced impacts compared to such a base case).

What do all these results tell us? Basically, we can conclude several things. A relative ranking of the varying forecasting and detruncation combinations was obtained, pointing to the fact that regression and projection is the best combination, while pickup and booking curve generated the lowest revenues for all seat optimizers tested. The revenue gains from a detruncation change are larger than those from a forecasting method change, although this is highly dependent on the choice of seat optimizer. In terms of the seat optimizers themselves, DAVN is a very robust method and performs rather well in all forecasting and detruncation combinations, while the other seat optimization algorithms are more sensitive to the forecasting and detruncation routines. Netbid is the most volatile; its performance appears to be quite sensitive to the choice of detruncation method. Booking curve detruncation appears to be a relatively ineffective method for Netbid, while projection detruncation causes it to perform best of all seat optimizers tested.

4.1.2.2 Actual leg loads

In addition to analyzing the overall system revenue to see how network revenues change for each airline given a certain combination of forecasting/detruncation method as compared with some base case set of methodologies as done in Section 4.1.2.1, further insight may be gained into why such results may occur by analyzing other outputs of the simulation. Therefore, this section attempts to look more into the inner workings of why the system revenue differences occur by examining the flight leg passenger loads. For example, system revenues may increase for two completely different reasons: (1) passenger loads generally increase on flight legs, or (2) passenger loads generally decrease but a better passenger mix is obtained (i.e., more higher-fare passengers and fewer lower fare passengers) on a large number of flight legs.

When examining the passenger load results, we wish to see how the different forecasting and detruncation combinations perform under the different seat optimization methods; and while system revenues are given for the network as a whole, passenger load analysis must be broken down either on a flight leg basis or a path basis. Figure 4.6 uses the flight leg approach in examining load changes by fare class. Shown are the *average* leg load changes[72] for Airline A for the particular fare class, expressed as the increase/decrease from a base case of pickup forecasting and booking curve detruncation at a demand factor of 1.0. Of each grouping of three columns in the graph, the leftmost group represents the Y-class loads, while the rightmost represents the Q-class loads; these are given for each of the seat optimizers tested.

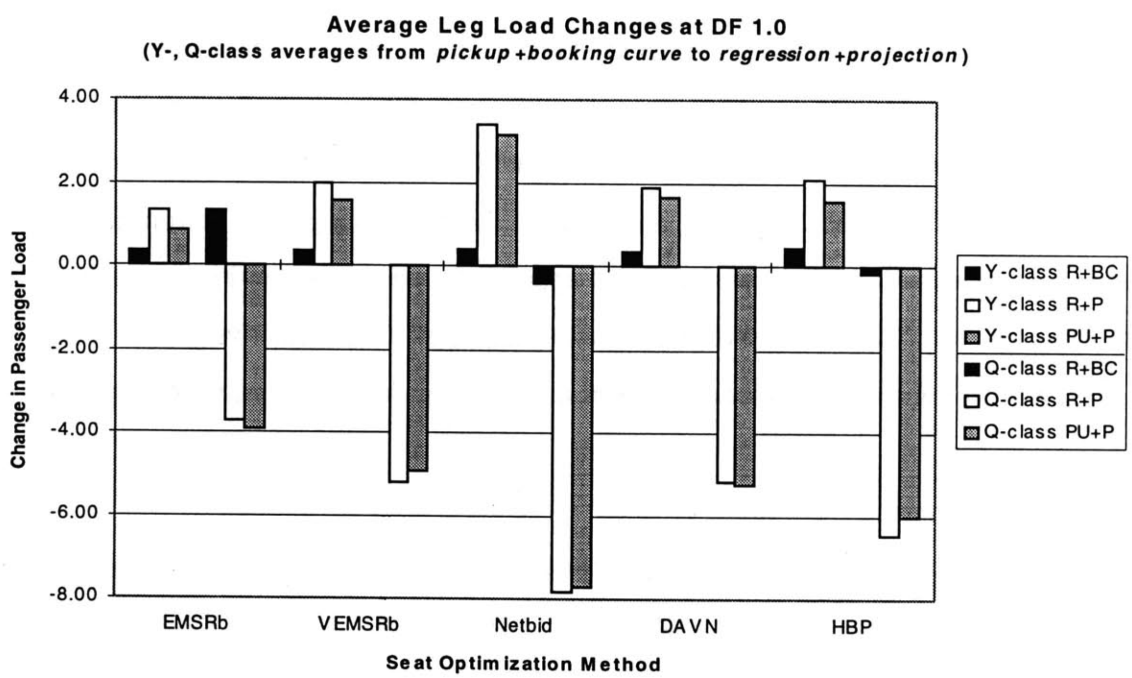

Figure 4.6: Y-,Q-class passenger load changes from forecasting/detruncation changes

[72] Averaged over all 12 flight legs of the network for the airline of interest.

Several conclusions can be deduced from the above graph. First, the passenger load changes for any particular forecasting/detruncation combination are again highly dependent on the choice of seat optimization routine (as was to be expected given the revenue results in Section 4.1.2.1). While the passenger load changes are almost uniform among seat optimizers when regression forecasting and booking curve detruncation is used, much larger changes occur when the detruncation method is changed to projection. In fact, Netbid has the largest magnitude of load change under projection detruncation, EMSRb has the smallest magnitude of change, and the other three methods have similar changes, of a magnitude in between that of EMSRb and Netbid. And for all seat optimizers, the trend is for Y-class bookings to increase when changing from the base case, while Q-class bookings decrease at a factor of about two times as large as the Y-class increases in bookings.

Second, the magnitude of the load changes is also dependent on the forecasting/detruncation inputs. A change from pickup to regression forecasting, using either of the two detruncation methods (illustrated by the dark-shaded blocks for booking curve detruncation and the *difference* between the white blocks and the medium-shaded blocks for projection detruncation in Figure 4.6), causes load changes on a much lower scale than a change in the detruncation method (represented as the medium-shaded blocks for pickup forecasting and the *difference* between the dark-shaded blocks and the white blocks for regression forecasting in Figure 4.6).

Y-class passenger load changes		***Seat Optimizer***				
		EMSRb	*VEMSRb*	*Netbid*	*DAVN*	*HBP*
DF 0.8	Switch forecast	0.29	0.23	0.28	0.20	0.27
	Switch detruncation	0.22	0.41	0.85	0.53	0.47
	Ratio (Detrunc/Fcst)	*0.8*	*1.7*	*3.0*	*2.6*	*1.7*
DF 1.0	Switch forecast	0.43	0.39	0.34	0.30	0.48
	Switch detruncation	0.92	1.61	3.07	1.60	1.61
	Ratio (Detrunc/Fcst)	*2.1*	*4.1*	*9.0*	*5.3*	*3.4*
DF 1.2	Switch forecast	0.58	0.30	0.12	0.26	0.56
	Switch detruncation	1.63	3.48	5.23	2.78	3.20
	Ratio (Detrunc/Fcst)	*2.8*	*11.8*	*44.3*	*10.8*	*5.7*

Table 4.3: Y-class passenger load changes for forecasting/detruncation changes

Tables 4.3 and 4.4 detail the average passenger load changes by seat optimization routine from switching the forecasting method (from pickup to regression) and switching the detruncation method (from booking curve to projection) while holding the other method constant under the three different demand factors.

Q-class passenger load changes		***Seat Optimizer***				
		EMSRb	*VEMSRb*	*Netbid*	*DAVN*	*HBP*
DF 0.8	Switch forecast	-0.22	-0.29	-0.36	-0.23	-0.41
	Switch detruncation	-1.04	-1.55	-2.36	-2.03	-1.86
	Ratio (Detrunc/Fcst)	*4.7*	*5.3*	*6.6*	*9.0*	*4.5*
DF 1.0	Switch forecast	0.77	-0.15	-0.27	0.04	-0.30
	Switch detruncation	-4.46	-5.07	-7.58	-5.21	-6.17
	Ratio (Detrunc/Fcst)	*-5.8*	*34.3*	*28.3*	*-135.9*	*20.9*
DF 1.2	Switch forecast	2.35	0.60	0.26	0.44	0.25
	Switch detruncation	-5.79	-9.03	-12.90	-7.68	-8.76
	Ratio (Detrunc/Fcst)	*-2.5*	*-15.1*	*-49.1*	*-17.4*	*-35.5*

Table 4.4: Q-class passenger load changes for forecasting/detruncation changes

As can be seen in the tables above, the opposite effects generally occur for Q- and Y-class. In the Y-class case, the average passenger load increase per flight leg varies from 0.1 to 0.6 passengers due to a change in the forecasting method to regression forecasting, regardless of the demand factor. However, a change to projection detruncation results in average Y-class passenger load increases of anywhere from 1.7 to more than 10 times greater (average passenger load increases per flight leg as high as 5.23 in the Netbid case at DF 1.2). In fact, only EMSRb at DF 0.8 has a detruncation/forecasting impact ratio less than 1.0, pointing to the fact that Y-class loads are much more affected by detruncation changes than forecasting changes. Additionally, as the demand factor is increased to 1.0 and 1.2, it can be seen that the passenger load changes for Netbid are of a much larger magnitude than those occurring under the other seat optimization routines, indicating the volatility of Netbid's passenger loads to the forecasting and detruncation methods.

In the Q-class case, the opposite effect occurs under demand factors 0.8 and 1.0; namely, passenger loads decrease when the forecasting method is changed to regression, but on a lower magnitude than the decreases experienced under the switch of the detruncation method to projection[73]. This trend is accentuated under DF 1.0; however, once the demand factor is increased to 1.2, we begin to see average Q-class leg loads increasing under all seat optimization methods as the forecasting method is changed to regression, and it is only the change to projection detruncation under which passenger loads decrease!

Third, the increase in network revenues resulting from implementation of projection detruncation seen in Section 4.1.2.1 appears to occur despite lower *total* loads, as Y-class loads generally increase with a factor of only about 1/2 to 1/3 of the Q-class load decrease in these cases (see Figure 4.6). This is also evident in Table 4.5 below, which compares the system average load factor (ALF) under the possible combinations of forecaster, detruncator, and seat optimizer.

	Forecasting+detruncation method	*EMSRb*	*VEMSRb*	*Netbid*	*DAVN*	*HBP*
DF 0.8	**Pickup+Booking Curve**	67.40	68.18	68.51	68.06	68.26
	Regression+Booking Curve	67.29	68.07	68.37	67.94	68.09
	Regression+Projection	66.70	67.11	67.06	66.51	66.95
	Pickup+Projection	66.82	67.25	67.20	66.61	67.19
DF 1.0	**Pickup+Booking Curve**	75.78	78.88	80.32	78.30	79.12
	Regression+Booking Curve	76.33	78.98	80.15	78.44	79.01
	Regression+Projection	73.80	75.62	75.48	74.19	75.06
	Pickup+Projection	73.72	75.73	75.51	74.09	75.31
DF 1.2	**Pickup+Booking Curve**	79.99	85.21	87.56	83.93	85.14
	Regression+Booking Curve	81.42	85.79	87.74	84.48	85.43
	Regression+Projection	77.38	80.04	79.36	77.62	78.93
	Pickup+Projection	76.92	79.80	78.96	77.24	78.92

Table 4.5: Average leg load factors for forecasting/detruncation changes

Note that in all cases, a change from booking curve detruncation to projection detruncation results in a decrease in load factor; however, for Netbid, this drop is most pronounced (from 80 to 75 under DF 1.0 and from 87 to 79 under DF 1.2). In fact, a comparison of the load factors under the different seat optimization routines for any case of forecasting, detruncation, and demand factor shows that Netbid is highest of the range when booking curve detruncation is used, while it is in line with the others when projection detruncation is used (VEMSRb consequently has the highest ALF when projection detruncation is used--the shaded blocks in Table 4.5

[73] Even at DF 0.8 the detruncation/forecasting ratio is always greater than 4.5.

represent the seat optimizer with the highest ALF). Therefore, Netbid's poor performance under booking curve detruncation can be seen primarily a result of loads being too high due to low-fare (i.e., Q-class) passengers (see Figure 4.7 below). In the opposite case, DAVN performs well under booking curve detruncation since it is still able to restrict the number of Q-class passengers accepted during the booking process. Again referring to Figure 4.7, it is evident that while Netbid performs comparably with DAVN under projection detruncation, under booking curve detruncation Netbid's Y-class loads are a bit lower and the Q-class loads are substantially higher, resulting in a worse passenger mix and lower overall revenues.

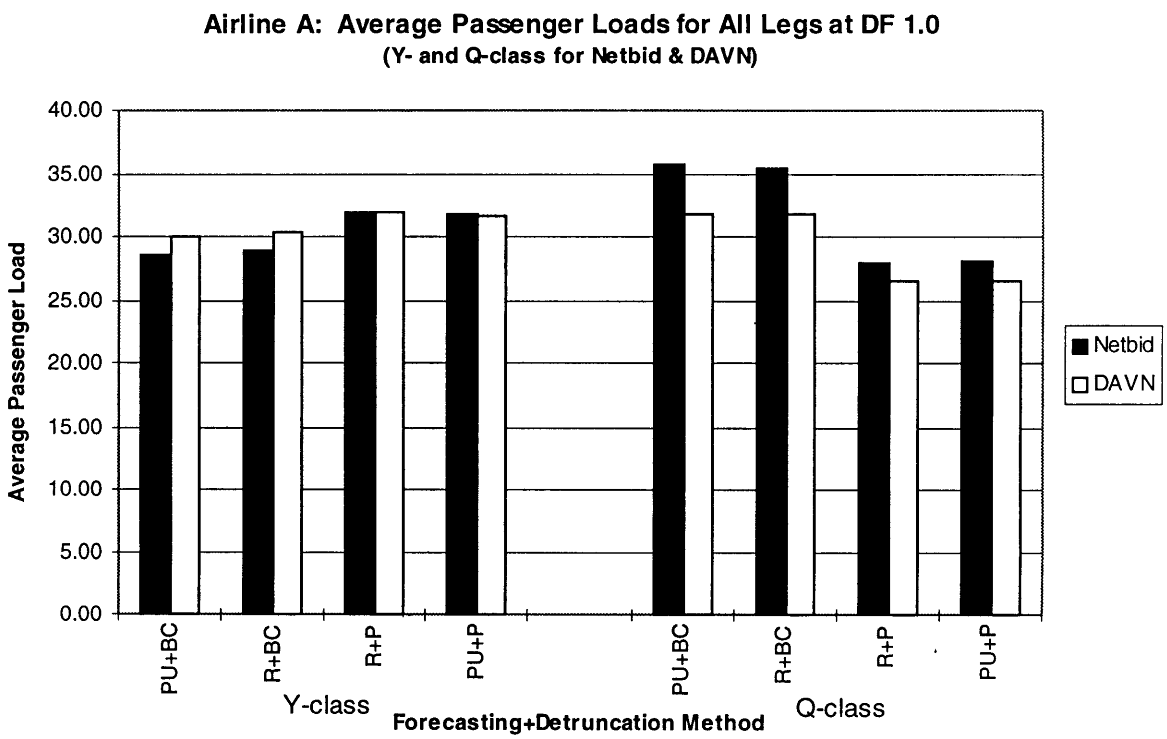

Figure 4.7: Average passenger loads for Netbid, DAVN

Lastly, as shown by the revenue results from Section 4.1.2.1, the forecasting/detruncation scenario producing the highest revenue for all seat optimizers occurred when projection detruncation was used, with the best case being the combination of regression forecasting and projection detruncation. Consequently, this is the scenario where the load changes have the largest magnitude. What can therefore be inferred is that implementation of projection detruncation results in demand predictions which cause a better allocation of seats to the individual fare classes than under booking curve detruncation, thereby resulting in higher Y-class loads whose revenue more than offsets the revenue loss from spilled Q-class customers. Put another way, booking curve detruncation appears to cause too many seats to be allowed to be filled by lower-fare customers which could have otherwise been taken by higher-fare passengers, especially in the Netbid case. But why does this discrepancy exist? One reason deals with that of Netbid's seat inventory control--it uses bid prices where other methods implement seat allocation limits. With seat allocation limits, excess low-fare bookings (due to higher demand than predicted) can be mitigated to some extent by the booking limits. With bid prices, no booking limits are in place; so all fare requests exceeding the bid price will be accepted (no matter how many requests are received)--this will continue until the next subsequent reoptimization of the booking limits. Hence the importance of frequent reoptimization with such seat inventory control.

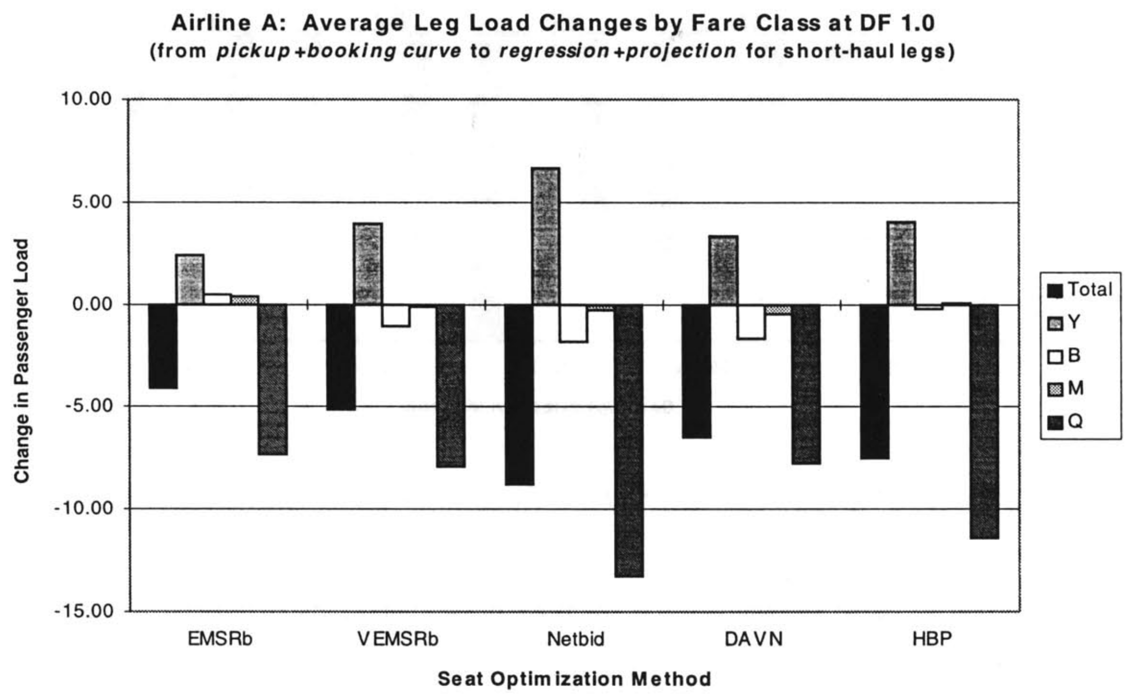

Figure 4.8: Leg load changes by fare class in short-haul markets at DF 1.0

As many different comparisons can be made among all the different input parameters, some selectivity must be used to illustrate the most important results. Based on the above analysis, the most interesting cases (and the most insightful) are those comparing the base case of pickup forecasting and booking curve detruncation to the "best case" of regression forecasting and projection detruncation. The next set of passenger load figures (Figures 4.8 and 4.9) therefore compares the load changes in going from the base case of pickup forecasting and booking curve detruncation to regression forecasting and projection detruncation only, broken down by fare class at the two higher demand factors (DF 1.0 and DF 1.2).

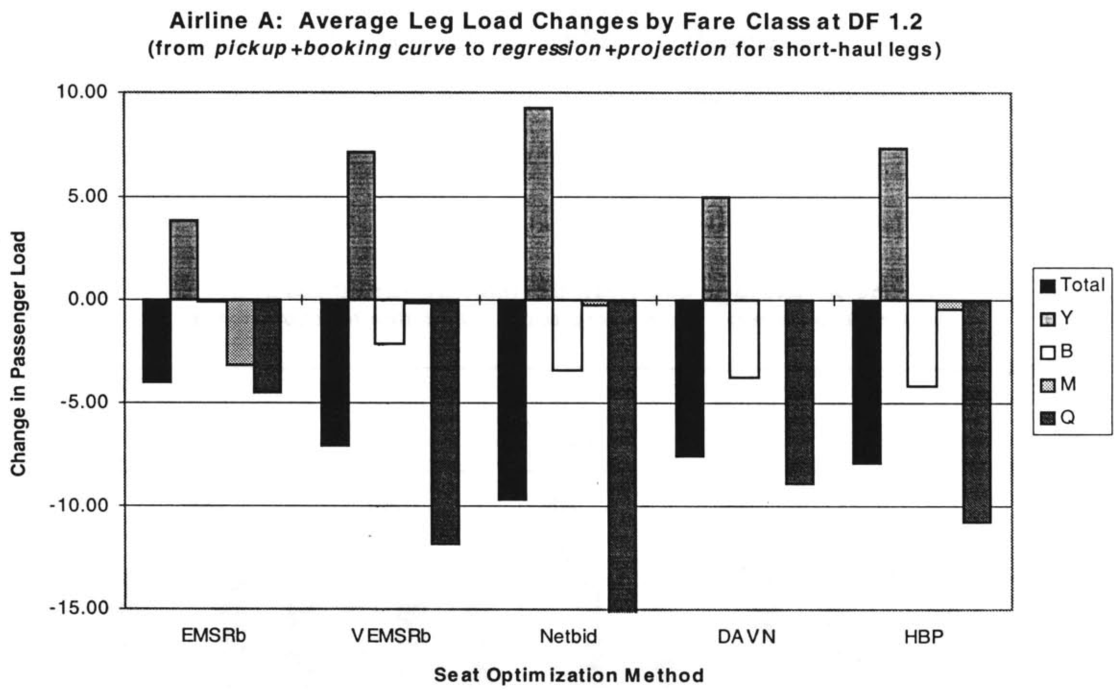

Figure 4.9: Leg load changes by fare class in short-haul markets at DF 1.2

The values shown in these graphs are the average passenger loads occurring on all short-haul (500-mile) legs (i.e., those with the highest demand). From this, two conclusions can be drawn. First, the B- and the M-class load changes are minuscule compared with those in Y- and Q-class for all seat optimization methods. This is mainly due to the fact that passenger bookings in these classes are much lower than bookings in Y- and Q-class, resulting from the so-called "bathtub effect[74]." Second, the load changes in the high demand factor cases are greatest, while those in the lower demand factor cases are smallest (the same Y-axis scale is used for all three cases to emphasize this point). This conforms with what is expected, as higher demand should produce higher changes in loads; therefore, although a better seat allocation is seen on all legs with the regression and projection combination, it is most pronounced on those with highest demand (i.e., DF 1.2). If this were not the case, we could conclude that the base case combination of methods was already doing an excellent job of obtaining the best passenger mix of fare classes.

The question may arise as to why only leg loads were analyzed in this section and path loads omitted. This is because in the PODS simulator, path (i.e., O-D) loads are less insightful than leg loads, as different markets are favored by different methods. That is, while a long-haul path may show the largest load increases for VEMSRb (refer to Section 2.3.1.2 for discussion), other shorter-haul paths may have larger increases for other seat optimization methods under the same set of algorithms. Therefore, the passenger load analysis by flight leg has provided good insight into some of the forces driving the revenue results obtained in the simulations. However, the next item of interest is that of the fare class closures, which themselves drive the passenger choice and path availability components and thereby the passenger loads and overall system revenues.

4.1.2.3 Fare class closures

In addition to the passenger loads resulting from the different forecasting, detruncation, and seat optimization combinations, another useful measure is the average time of fare class closure. For any given ODF, the average time frame in which the fare class closed can be determined (i.e., the point in time after which no additional bookings can be accepted). Figure 4.10 shows these fare class closure times for a single-leg path (from the hub of Airline A to a spoke city) at DF 1.0. The path analyzed was a short-haul, 500-mile path, with rather high demand.

The darker the shaded region, the more fare classes are open (i.e., all fare classes are available in the black regions, while only Y-class is open in the white regions). Again, only the

[74] The "bathtub effect" is the PODS phenomenon that Y- and Q-class bookings are of relatively larger magnitude than their B- and M-class counterparts. This is primarily due to the perceived cost of the fares to business/leisure passengers based on the relative disutilities of the different restrictions.

regression/projection combination was compared against the pickup/booking curve method. For the Q-class fares, it can be seen that under regression/projection, the fare classes are closed for all methods which employ a virtual fare class scheme (i.e., VEMSRb, DAVN, and HBP), and even under pickup/booking curve only a couple of virtual bucketing methods result in Q-class fares being open, albeit minimally). For M-class, virtually the same situation exists (in both cases because the Q- and M-class fares occupy the lowest virtual bucket, for which very few seats are allocated given such a medium-demand scenario). In the non-virtual class schemes (i.e., EMSRb and Netbid), we see that while Q- and M-class fares are open for some amount of time, they close much sooner in the regression/projection case than in the base case of pickup/booking curve.

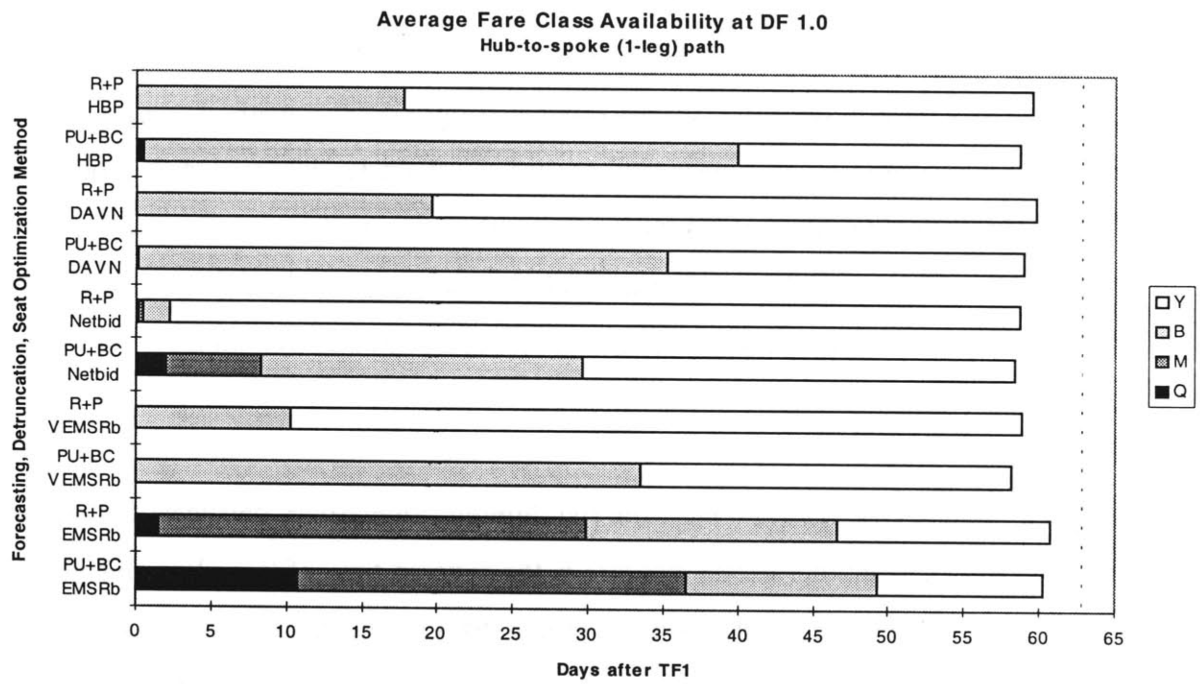

Figure 4.10: Average fare class closure times for a hub-to-spoke path at DF 1.0

It is rather in B-class and Y-class where dramatic changes become evident under different forecasting routines. Under regression/projection, the B-class fares always close earlier than under pickup/booking curve; this difference is substantial for all seat optimizations except EMSRb. Although passenger load analysis has indicated that B-class loads are much lower than their Y-class counterparts, the day (or time frame) of B-class closure is important in that once it closes, only Y-class fares are available. Examination of Figure 4.10 shows that in all cases of regression/projection except for EMSRb, only Y-class fares are available after day 20 (the fourth time frame), and even as early as day 3 (the second time frame) in the Netbid case. As for the Y-class fares, most of them close near the end of the booking process regardless of forecast method/detruncation method/seat optimizer choice, meaning that the Y-class fares were usually available for most of the process. Bringing this into the context of relative revenue performance, we can see that Netbid's high revenues under regression/projection can probably be attributed to the fact that the Q- and M-class fares are never available on this path and that only the Y-class fares are open to passengers for most of the booking process, due to higher bid prices for the lower fare classes. In all other forecasting/detruncation and seat optimization combinations for the same path, passengers often have Q- and M-class available, with the option of B-class being available for a much more significant portion of the booking process.

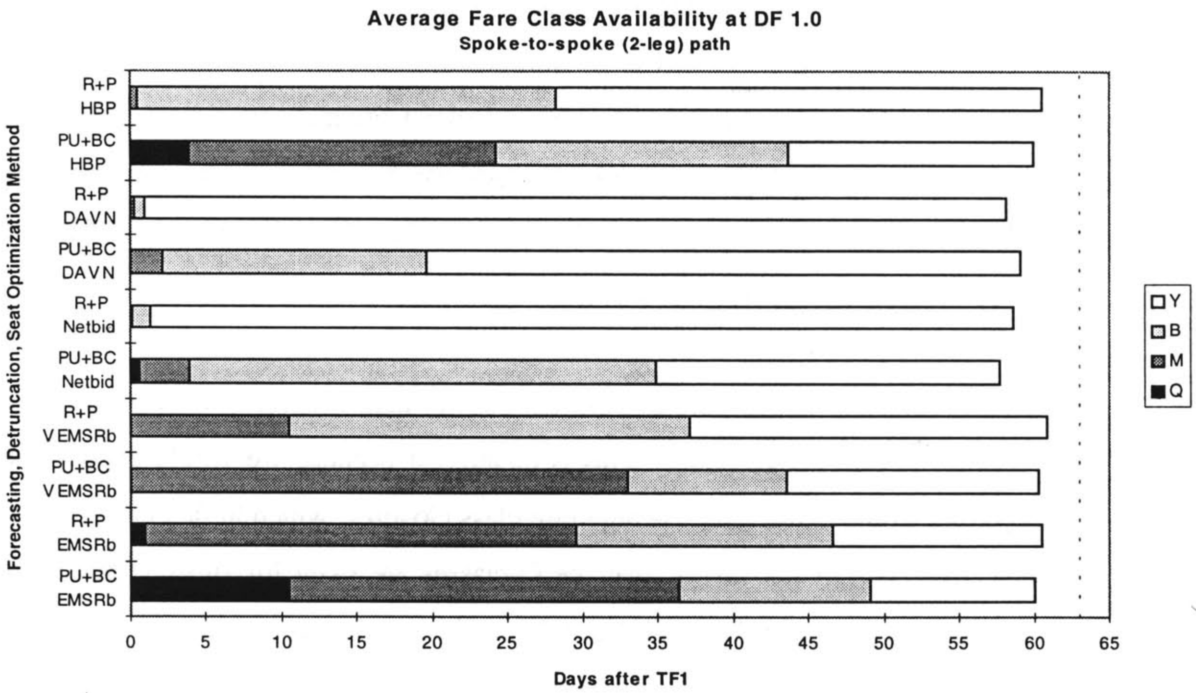

Figure 4.11: Average fare class closure times for a spoke-to-spoke path at DF 1.0

Figure 4.11 presents similar data as in Figure 4.10, but this time for a path traversing two legs (i.e., connecting two spoke cities through a hub), rather than a single-leg one. Here the path analyzed was a 1000-mile spoke city to spoke city path with moderately high demand, although not as high as in the first case, as the PODS system inputs were defined to have local demand higher than connecting demand (see Section 3.2).

Many of the same trends are evident for the spoke-to-spoke path as for the hub-to-spoke path, although here we see that for seat optimizers using a virtual bucketing scheme, M-class and even Q-class fares are available at some time frames (this is to be expected since in the single-leg case from before, the Q-class fares are generally lower in price than the Q-class fares present in these two-leg O-D markets, and are therefore assigned to the lowest virtual bucket, which ends up receiving few or no seat allocations). Again, the difference between Q-, M-, and B-class closure is very pronounced when changing from pickup/booking curve to regression/projection, especially in the Netbid case. Also similar to the single-leg case, Netbid's good performance probably results from only having the Q- and M-class fares completely closed, with only Y-class fares available for most of the booking period despite the lower demand situation as compared with the hub-to-spoke path (DAVN's equally high revenue in this case can be accounted for similarly). In fact, part of DAVN's robustness can be seen in that it has the earliest fare class closures of all the seat optimization methods in the pickup/booking curve case.

Finally, Figure 4.12 also presents similar data as Figure 4.10, but here it is for a demand factor of 0.8 (on the same one-leg hub-to-spoke path). This is of interest because we would expect that such lower demand cases would cause later average fare class closures, which in fact can be seen. One main difference from the DF 1.0 case is that the Q-classes are open for substantially more time than in the DF 1.0 case, meaning that in the low-demand scenarios, more seats are allocated to this lowest path fare class under the base case forecaster and detruncator. For all cases of O-D seat optimizer, it can still be seen that the regression/projection combination advances the fare class closure time substantially, causing Q- and M-class fares to close rather early in the booking process.

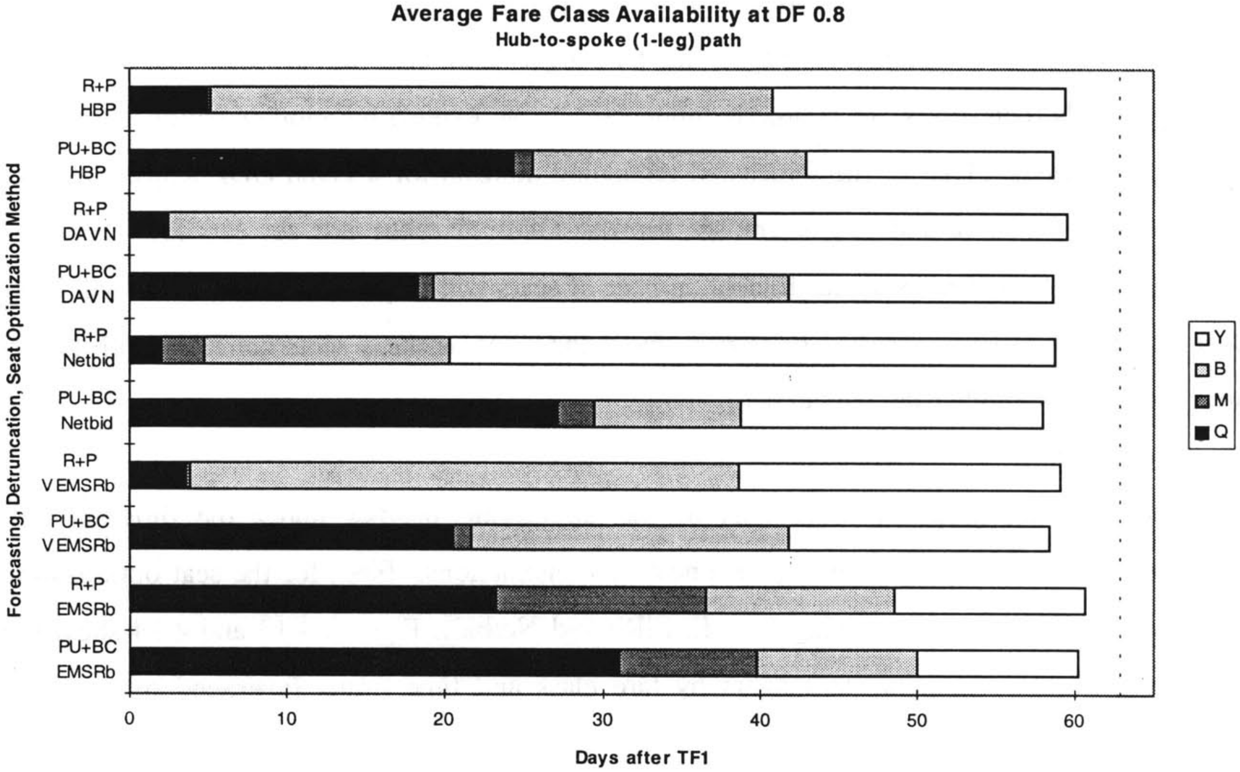

Figure 4.12: Average fare class closure times for a hub-to-spoke path at DF 0.8

Basically, what can be concluded from the fare class closure analysis is that employing the combination of regression/projection causes the Q-, M-, and B-class fares to close earlier than in the pickup/booking curve case. And under a medium-demand (DF 1.0) scenario, there is evidently enough extra Y-class demand to offset the loss of lower fare class demand by the earlier closing of these lower fare classes, resulting in higher Y-class loads and higher system revenues.

4.1.2.4 Forecasted demands

The analysis in the previous three sections has shown that revenues under Netbid are quite sensitive to the forecasting and detruncation inputs. Therefore, another item of interest is that of the driving force behind what influences the fare class closures, passenger loads, and system revenues; namely, forecasted remaining passengers. Analysis of these values can also provide insight into such phenomena as why booking curve detruncation performs poorly with Netbid.

Although our primary concern is with overall system revenues, these revenue values are driven by passenger loads, as explained above in the passenger load analysis section. However, the passenger loads themselves are indirectly influenced (to an extent) by another factor--forecasted remaining demands. That is, the forecasted remaining demand for a given ODF helps determine what booking limits or bid prices are set for the ODF; if these are set too low, high-fare passengers may be spilled as an insufficient number of seats will be protected, and loads will likely decrease (in lower-demand cases) or remain about the same as these seats are filled with lower-fare passengers (in higher-demand cases).

Presented next are graphical depictions of the forecasting profiles under the different seat optimization methods and forecasting/detruncation combinations. First, for the seat optimization methods not using virtual bucketing (i.e., EMSRb and Netbid), Figures 4.13 and 4.14 show the average forecasted remaining passengers by fare class and time frame (averaged over the 20 trials).

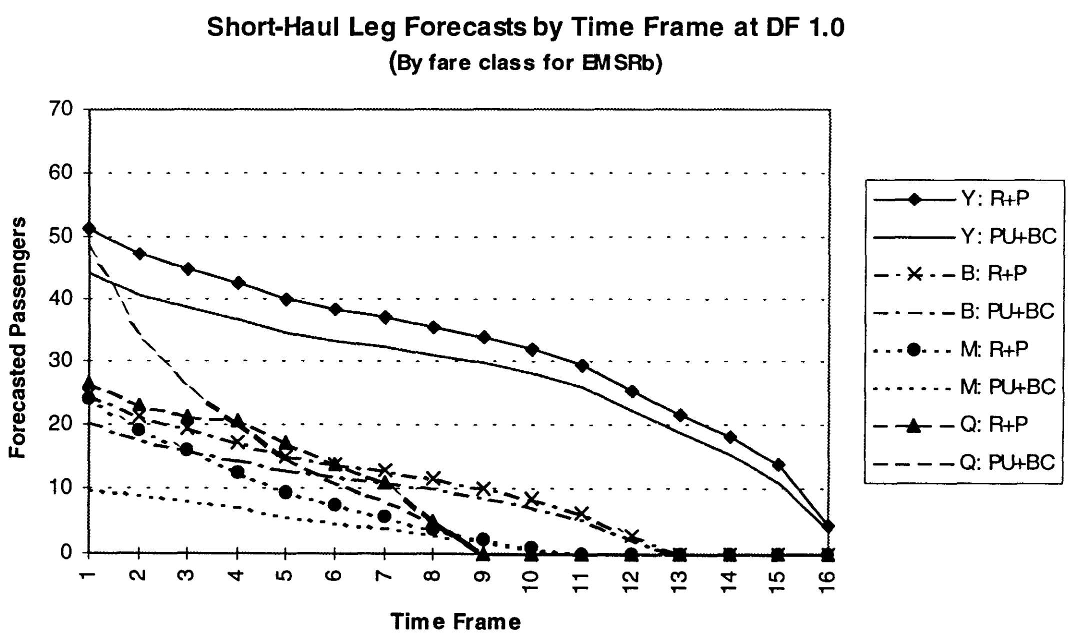

Figure 4.13: Short-haul leg forecasts by time frame for EMSRb

As can be seen, the relative magnitudes of the B- and M-class forecasts and passenger bookings are much lower than their Y- and Q-class counterparts; hence, the differences obtained by initiating a forecasting or detruncation method change are much smaller. It should also be noted that while forecasts for EMSRb are obtained on a flight leg basis (as EMSRb uses leg-based control), Netbid uses path control in its seat optimization algorithm. Therefore, to make valid comparisons, all paths traversing a leg were summed to gain an equivalent leg forecast, which can subsequently be compared with the leg forecast made in the EMSRb case. In the examples that follow, a leg from the hub of Airline A to spoke city A was used, for which all paths traversing this leg were summed in the Netbid case.

Figure 4.13 illustrates the forecasting profiles for a representative short-haul, high-demand leg when EMSRb is the choice of seat optimization method, while Figure 4.14 details the same data but for Netbid as the seat optimization choice.

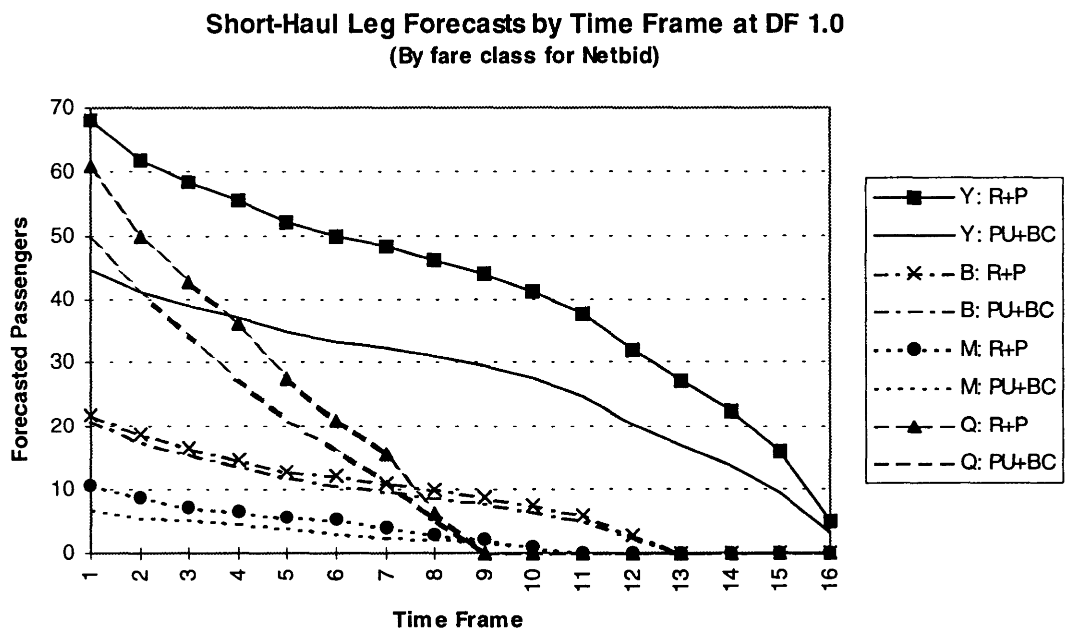

Figure 4.14: Short-haul path forecasts by time frame for Netbid

When the combination of regression forecasting and projection detruncation is used, the forecasts tend to always be higher in Y-class (as well as B- and M-class), but how much higher these forecasts are differs by seat optimization method. With Netbid, implementation of regression/projection causes Y-class forecasts to be about 50% higher than the base case on average, while under EMSRb the same change only increases them by about 10-15% over the base case of pickup forecasting/booking curve detruncation. On the other hand, the same forecasting/detruncation change to regression/projection causes Q-class forecasted bookings to go up by about 20-30% in the Netbid case, while in the EMSRb case, this change actually causes lower forecasts for the first few time frames than in the base case! This could be due to the fact that the pickup/booking curve combination just overestimated initial Q-class demand in the EMSRb case--from which it took several time frames for forecasts and bookings to equilibrate.

In essence, it appears that Netbid is performing better in terms of revenue because projection detruncation causes higher forecasts in the higher fare classes; and this demand ends up materializing. If it did not materialize, we would expect that Netbid would perform *worse* as seats in higher fare classes would be protected and then go empty, while at the same time lower fare class passengers would be spilled. What can therefore be inferred is that Netbid performs worse than EMSRb when the forecasts are too low, since lower forecasts induce lower bid prices on the network paths, as passenger displacement costs are lower. On the other hand, EMSRb technology sets booking limits for the fare classes based on the forecasts. And while a lower forecasted demand will cause lower fare class limits, these limits cannot be exceeded by the demand, whereas with Netbid, there is no capacity limitation on the bid prices and any demand whose fare is greater than this bid price will be accommodated. Furthermore, as we saw that Netbid's fare classes were closed more often, which requires more need for detruncation, thereby giving detruncation a larger overall impact.

A hypothetical example helps illustrate this point. Assume that we have one flight leg with an aircraft capacity of 10, two fare classes (call them Y and Q), each with demand of 10, and fares of \$400 in the Y-class and \$100 in the Q-class. In the case of low forecasted demands (say 5 in each class), the EMSRb seat optimization routine should reserve 5 seats for the Y-class and therefore allocate the remaining 5 seats to the Q-class; and let us assume that the Netbid routine sets the bid price of the leg to be \$90. Using EMSRb, we would experience 5 Q-class bookings followed by 5

Y-class bookings with resultant revenues of $2,500, while with Netbid all 10 Q-class passengers would be accepted since the $100 fare is greater than the bid price and would thereby displace the Y-class customers, resulting in total revenues of $1,000. Therefore, it can be seen that Netbid's lack of capacity restrictions can have adverse consequences if the forecasts are inaccurate and low, and while although not beneficial for EMSRb either, the adverse effects are mitigated to the extent that there is a limiting capacity in place.

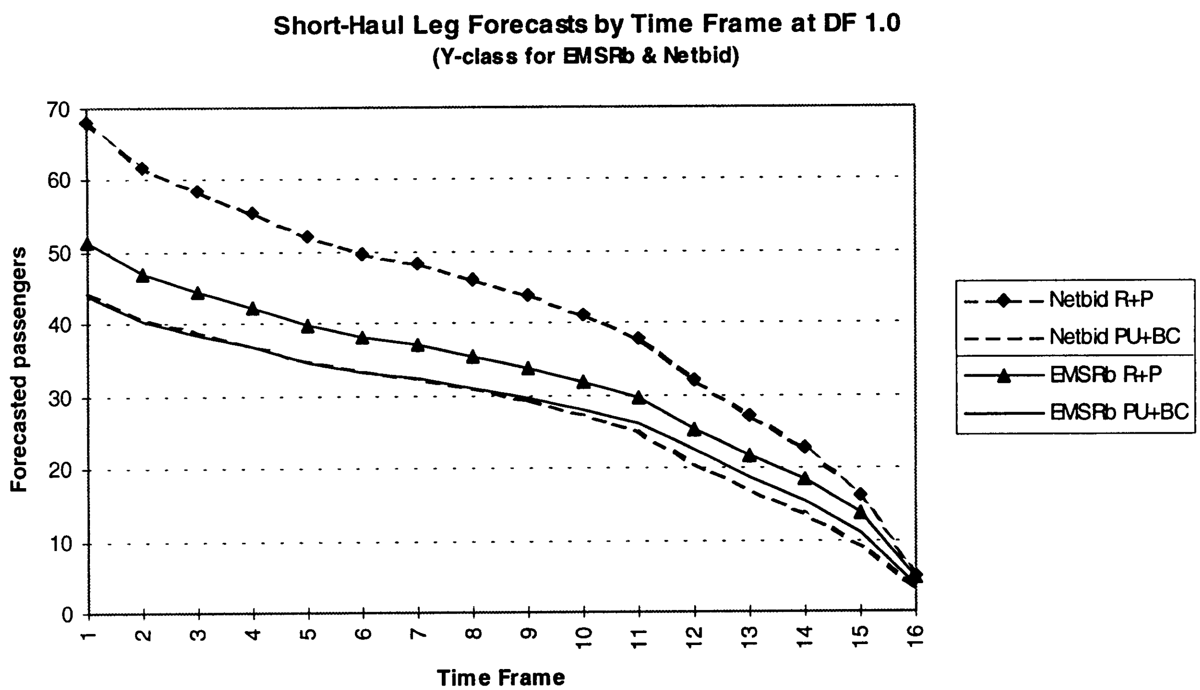

Figure 4.15: Forecasted Y-class demands by time frame for EMSRb, Netbid

Furthermore, Figure 4.15 above makes a simple comparison between EMSRb and Netbid and shows for just the Y-class how the forecasted remaining demands perform under these two seat optimization methods with the two combinations of forecaster and detruncator. Under the pickup/booking curve combination, both EMSRb and Netbid have similar forecasts, while under regression/projection the Netbid forecast increases substantially more than the EMSRb forecast does, lending support to the theory that Netbid performs best under a forecasting/detruncation combination that makes higher forecasted remaining demands in the highest fare classes.

Next, the seat optimization methods using virtual fare classes (i.e., VEMSRb, DAVN, and HBP) were compared. Such comparisons are less insightful, however, since a Y-class fare in a short-haul market can map into the same virtual bucket as an M- or even a Q-class fare in a long-haul market, given the various fare characteristics of the different O-D markets. Hence, conclusions about trends among the different fare classes are difficult to infer, although comparisons can still be made between the different scenarios of forecasting and detruncation method combinations. Shown below in Figure 4.16 is such a comparison--the percentage increases for the forecasted loads by virtual fare bucket, from which it can be seen that on average, forecasts made under the regression/projection combination are approximately 25-30% higher for each of the three seat optimization methods employing virtual fare buckets.

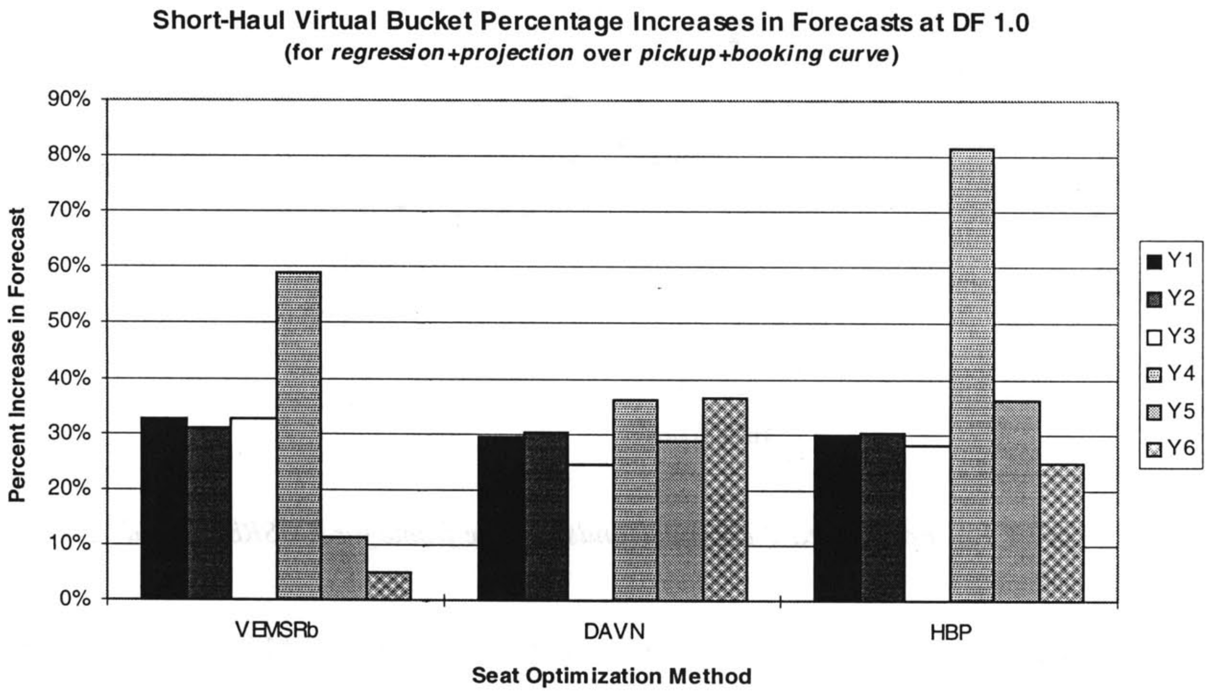

Figure 4.16: Virtual class percentage increases for a short-haul leg forecast

Based on the above discussion, what can be concluded? Because all the seat optimization algorithms implemented use top-down seat protection, the forecasts with the largest impacts on path (and system) revenues will be those in the highest fare class for a particular ODF (Y-class for non-virtual class seat optimization, the virtual class to which the Y-class fare maps to for virtual class seat optimization), followed by those in the next lowest fare class, with the lowest fare class forecasts only determining seat protection for the lowest fare class, if any seats remain which can be protected. So the best performance combination will be the one that most accurately predicts the demand in the highest fare class. We have already postulated that the detruncation scheme has a far more influential impact on our demand forecasts; and based on the empirical evidence presented above, we can conclude that among our detruncation options, the booking curve method just produces too low of a forecast, affecting any seat optimizer that is sensitive to underpredictions of demand.

In conclusion, analysis of the actual loads by time frame, the fare class closure results, and the forecasted passenger loads and by departure (both on a leg basis and a path basis), has provided some explanation about the reasons why different system revenue values under different combinations of seat optimizer, forecasting routine, and detruncation method were occurring (see Section 4.1.2.1). From the load analysis combined with our revenue results, we have determined that the choice of detruncation method has a much larger impact on both loads and system revenues. Furthermore, as projection detruncation is responsible for much of the load changes seen under the different methods; the combination of regression forecasting and projection detruncation produces the largest difference with respect to the base case of pickup forecasting and booking curve detruncation; subsequent analysis will therefore be concentrated upon those changes.

Additionally, in terms of revenue performance, we have seen that in general, revenues increase when projection detruncation is employed, while passenger loads decrease. However, this decrease is a result of Q-class loads decreasing more rapidly than a corresponding increase in Y-class loads (it is this increase that is mainly responsible for the large revenue gain). Furthermore, these loads are directly influenced by the time frames in which fare class closure occurs--the method combinations with higher Y-class loads (and hence, revenues) correlate with those methods in which the lower fare classes close earlier in the booking process.

4.2 Impacts of competitive scenarios

Section 4.1 has provided us with a good analysis of the results occurring when Airline A changes its forecasting or detruncation method in combination with its seat optimization routine. However, in all cases analyzed, Airline B was held constant using a fixed base case of EMSRb seat optimization technology with pickup forecasting and booking curve detruncation. Therefore, this section will allow more freedom on the input level by comparing the relative performances of forecasting and detruncation routines in competitive scenarios--that is, those in which both airlines have some degree of flexibility in their choice of forecaster and detruncator. Section 4.2.1 will examine cases where both airlines are constrained to using the same seat optimization method. In essence, what will be obtained is a category of results which measure the relative performances of the forecasting and detruncation methods for head-to-head competition under each of the five specific seat optimizers. Section 4.2.2 will then analyze base cases where a seat optimization routine other than EMSRb is used by the competitor--that is, full flexibility in forecaster, detruncator, and seat optimizer choice is given to both airlines. Finally, Section 4.2.3 will present a synopsis of the best combinations of results. Total system revenue will again be used as a metric for comparing the performance, as the trends for passenger loads, forecasts, and fare class closures are similar to what was discovered and presented in Section 4.1.

4.2.1 Seat optimization method changes by both competitors

In the scenarios analyzed in Section 4.1, Airline B used a base case of EMSRb with pickup forecasting and booking curve detruncation, and therefore Airline A's forecasting, detruncation, and seat optimization methods were always competing against this scenario, which was seen to be the least revenue-beneficial of all of the combinations under all demand factors. Although such an analysis was important in recognizing and isolating the relative trends among the methods, the revenue results and gains for Airline A were being inflated to some degree by the fact that the competition was always subjected to an "inferior" combination of methods. Therefore, Table 4.6 below details the percent increase in revenue over the base case of EMSRb with pickup and booking curve[75] at a demand factor of 1.0, when both airlines use the seat optimization method

[75] The percent increase in revenue due solely to the forecasting/detruncation method can still be computed simply by taking the desired value in the table and subtracting the value in the first row (e.g., for the 2.47% revenue

listed in the top of the table (Airline B will still be using the base case of pickup forecasting and booking curve detruncation). Table 4.7 presents the same information for a demand factor of 0.8.

Airline A *Forecast/Detruncation*	*EMSRb*	*VEMSRb*	*Netbid*	*DAVN*	*HBP*
Pickup/Booking Curve	-----	0.50%	0.15%	0.80%	0.60%
Regression/Booking Curve	0.40%	0.93%	0.65%	1.23%	0.99%
Regression/Projection	1.12%	1.48%	2.47%	1.15%	1.68%
Pickup/Projection	0.55%	1.05%	2.18%	0.88%	1.28%

Table 4.6: Percent revenue increases for identical seat optimization methods at DF 1.0

Airline A *Forecast/Detruncation*	*EMSRb*	*VEMSRb*	*Netbid*	*DAVN*	*HBP*
Pickup/Booking Curve	-----	0.09%	-0.15%	0.14%	0.08%
Regression/Booking Curve	0.35%	0.41%	0.34%	0.44%	0.42%
Regression/Projection	0.47%	0.60%	0.93%	0.53%	0.63%
Pickup/Projection	0.15%	0.36%	0.69%	0.37%	0.40%

Table 4.7: Percent revenue increases for identical seat optimization methods at DF 0.8

The tables above highlight several important points. First, the same relative trends seen in Section 4.1 still hold here. That is, Netbid performs best when projection detruncation is used, while DAVN is the best seat optimizer under booking curve detruncation (the shaded blocks represent the highest revenue value for the forecasting/detruncation combination). Furthermore, HBP can be considered a very "consistent" seat optimizer, in that it outperforms DAVN in the projection detruncation cases by more than it is outperformed in the booking curve detruncation ones. Second, it is evident that the revenue increases were due in some part to the better seat optimization routines and not completely to the fact that the competition was using the least beneficial combination of methods. This can be more clearly illustrated by Figure 4.17.

---

increase with Netbid/regression/projection, 0.15% is due to Netbid alone, while 2.47 - 0.15 = 2.32% is due to the better forecasting and detruncation combination).

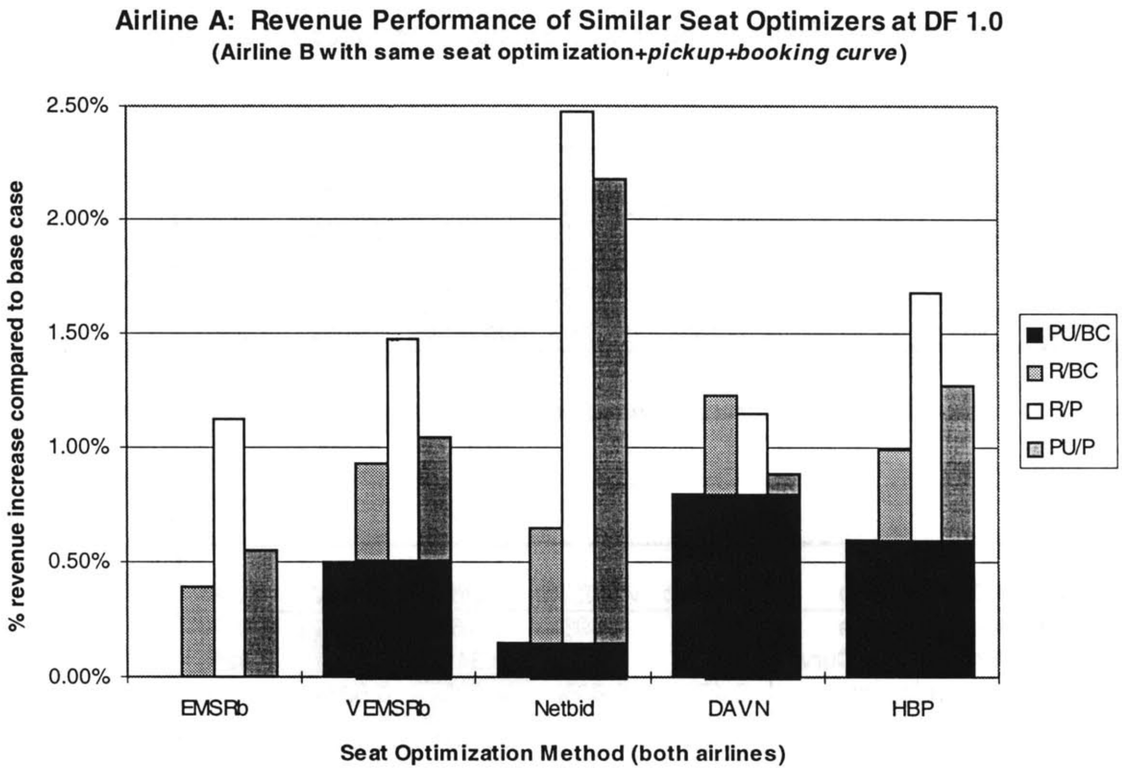

Figure 4.17: Revenue performance of similar seat optimizers for both competitors

Upon examination of the graph, it can be seen that while the percentage increases in revenue due to both competitors changing their seat optimization routine only (illustrated by the black columns) are much smaller than the gains obtained by Airline A when it changes its forecasting/detruncation inputs (illustrated by the areas above the black columns). Additionally, while the percentage increase is almost uniform by seat optimization method in changing to regression/booking curve (depicted by the darkly-shaded blocks), one can see that when projection detruncation is implemented, the percentage gains under Netbid are much higher than average; those under DAVN are much lower (represented by the white and lightly-shaded blocks). This reiterates the fact about the robustness of DAVN and the volatility of Netbid to forecasting/detruncation inputs, even when the competition is also using an equivalent seat optimization method.

4.2.2 Alternative yield management base cases

While Section 4.1 detailed the analysis of changes by Airline A against a fixed base case of methodologies used by Airline B and Section 4.2.1 described cases in which both airlines used the same seat optimization method (but Airline B still had a fixed base case of pickup forecasting and booking curve detruncation), this section expands into the cases where all combinations of forecasting, detruncation, and seat optimization are possible for both airlines. What this provides is a representation of the simulated characteristics of various scenarios of head-to-head competition. It is useful in that one can deduce the best set of inputs to use for a particular airline, given the methodological inputs of the competing airline; in addition to a set of results indicating the best performance cases to use overall, given that the competition is using one or more characteristic methodologies.

There are numerous possible cases of different combinations of forecaster, detruncator, and seat optimizer which can be compared for the two airlines; however, these cases can be simplified into five major categories--one for each of the different seat optimization methods used by the competition (Airline B). Within each of these five categories, there are still numerous head-to-head cases which occur (depending on the forecasting and detruncation method choices for both airlines as well as the seat optimization algorithm chosen by Airline A). Hence, the following discussion will provide some insight into the trends observed under different scenarios when it is known which seat optimizer is being used by the competition. From this, a preferred set of inputs can be inferred for Airline A based on the PODS simulation results.

For the scenario where the competing airline chooses VEMSRb seat optimization technology, revenues are still only moderately affected by forecasting/detruncation method changes, although to a higher degree than with EMSRb technology. Similar to before, if the competition will be using VEMSRb, Airline A should again choose Netbid if they will be using projection detruncation, and DAVN if booking curve detruncation will be used.

As has been shown earlier, Netbid's success is highly dependent on the choice of the combination of forecasting/detruncation method; changing the detruncation method to projection for Airline B has the same expected effect if it is implementing Netbid. Such cases result in moderate decreases in revenue for Airline A (except when Airline A also uses projection detruncation, which then causes slight revenue increases for both airlines). In addition to the usual case of Netbid performing best when projection detruncation is used by Airline A and DAVN otherwise, we see that HBP is also a good candidate for competing against a carrier using Netbid with projection detruncation.

The DAVN seat optimizer has been shown to be a highly robust choice; when it is implemented by the competition, revenue gains for Airline A are either sharply decreased or often negative, depending on its combination of forecaster/detruncator. In the cases where revenue gains do occur, the forecast/detruncation method changes produce gains which are generally less than 0.5%, mainly because of DAVN's good overall performance and the subsequent lack of gains due to "competitor feedback" (as discussed later in Section 4.4.2). As in the Netbid case, HBP is also an effective tool for competing against an airline using DAVN.

An interesting note should be made about the HBP seat optimization method--while it never ranks first as a method choice under any situation (only DAVN and Netbid do), it does however perform very consistently. That is, while DAVN performs best under booking curve detruncation and Netbid under projection detruncation, HBP is not that far behind either one, and it performs moderately better than either of those seat optimization methods under the inferior choice of the two detruncation methods (i.e., booking curve detruncation).

The final step in the analysis is that of fully competitive scenarios, where both Airlines A and B have the ability to vary any of their forecasting, detruncation, or seat optimization methodologies. Illustrated in Tables 4.8 and 4.9 are just two examples of the myriad of different combinations possible between the two airlines--the examples shown are those for which the competitor (Airline B) is using one of the "best" combinations of methodologies. The first, in Table 4.8, shows the revenue percentage increases (or decreases) from a base case of DAVN with pickup forecasting and booking curve detruncation, for both Airlines A and B in the case where Airline B is using DAVN with regression forecasting and projection detruncation.

Forecast/Detruncation		Airline A Seat Optimizer EMSRb	VEMSRb	Netbid	DAVN	HBP
Airline A	PU+BC	-2.13%	-0.74%	-1.22%	-----	-0.40%
	R+BC	-1.77%	-0.29%	-0.79%	0.42%	-0.08%
	R+P	-1.21%	0.28%	0.56%	0.23%	0.53%
	PU+P	-1.71%	-0.12%	0.29%	-0.07%	0.17%
Airline B	PU+BC	1.80%	0.34%	1.07%	0.00%	0.27%
	R+BC	1.83%	0.26%	1.03%	-0.03%	0.26%
	R+P	1.14%	-0.33%	-0.20%	-0.27%	-0.54%
	PU+P	1.29%	-0.25%	-0.19%	-0.25%	-0.48%

Table 4.8: Revenue percentage gains vs. Airline B using DAVN/regression/projection

We have seen that DAVN performs rather well in all cases, especially those with regression forecasting and projection detruncation. Therefore, when the competition is using such a yield management combination, it can be seen that the revenue gains for Airline A are much lower in magnitude, especially when the inferior forecasting and detruncation methods are used (positive revenue gains over the base case are shaded). For example, we can see that no positive revenue gains for Airline A can be achieved with the base case of pickup forecasting and booking curve detruncation, and for regression forecasting with booking curve detruncation the only positive revenue increases are obtained with DAVN. It is only with projection detruncation (and more specifically the regression/projection combination) that positive revenues can be achieved for the other seat optimizers such as Netbid and HBP. In Table 4.9 we again see the revenue percentage increases (or decreases) from a base case of Netbid with pickup forecasting and booking curve detruncation, for both Airlines A and B in the case where Airline B is using Netbid with regression forecasting and projection detruncation (as this was the "best" case of methodologies).

Forecast/Detruncation		Airline A Seat Optimizer EMSRb	VEMSRb	Netbid	DAVN	HBP
Airline A	PU+BC	-2.15%	-0.66%	-1.22%	-----	-1.20%
	R+BC	-1.81%	-0.24%	-0.81%	0.41%	-1.27%
	R+P	-1.27%	0.28%	0.54%	0.25%	0.51%
	PU+P	-1.74%	-0.11%	0.29%	-0.05%	0.25%
Airline B	PU+BC	1.74%	0.23%	0.95%	0.00%	0.67%
	R+BC	1.78%	0.15%	0.91%	-0.05%	0.87%
	R+P	1.05%	-0.41%	-0.29%	-0.27%	-0.54%
	PU+P	1.18%	-0.33%	-0.29%	-0.23%	-0.40%

Table 4.9: Revenue percentage gains vs. Airline B using Netbid/regression/projection

Here similar trends to Table 4.8 are seen, in which negative revenue gains are obtained by Airline A under booking curve detruncation, while small revenue gains can be achieved with projection detruncation. Basically, this synopsis illustrates that when the competition implements one of the "best" yield management combinations, positive revenue gains can be achieved mainly by improving the detruncation method on the different O-D seat optimizers, where the magnitude of these gains varies by seat optimization method (note that no case of forecaster/detruncator yields positive revenue gains for Airline A if it uses EMSRb in such a case).

4.2.3 Best forecasting/detruncation method for seat optimization methods

Another approach to the analysis of the relative performance of the seat optimization routines can also be taken in order to answer the question of how large the revenue gains are for each method if the "best" case of forecaster and detruncator is always being used. Figure 4.18 answers this exact question at three different demand factors by showing the percent revenue gain achieved by Airline A over a base case of EMSRb with regression forecasting and projection detruncation.

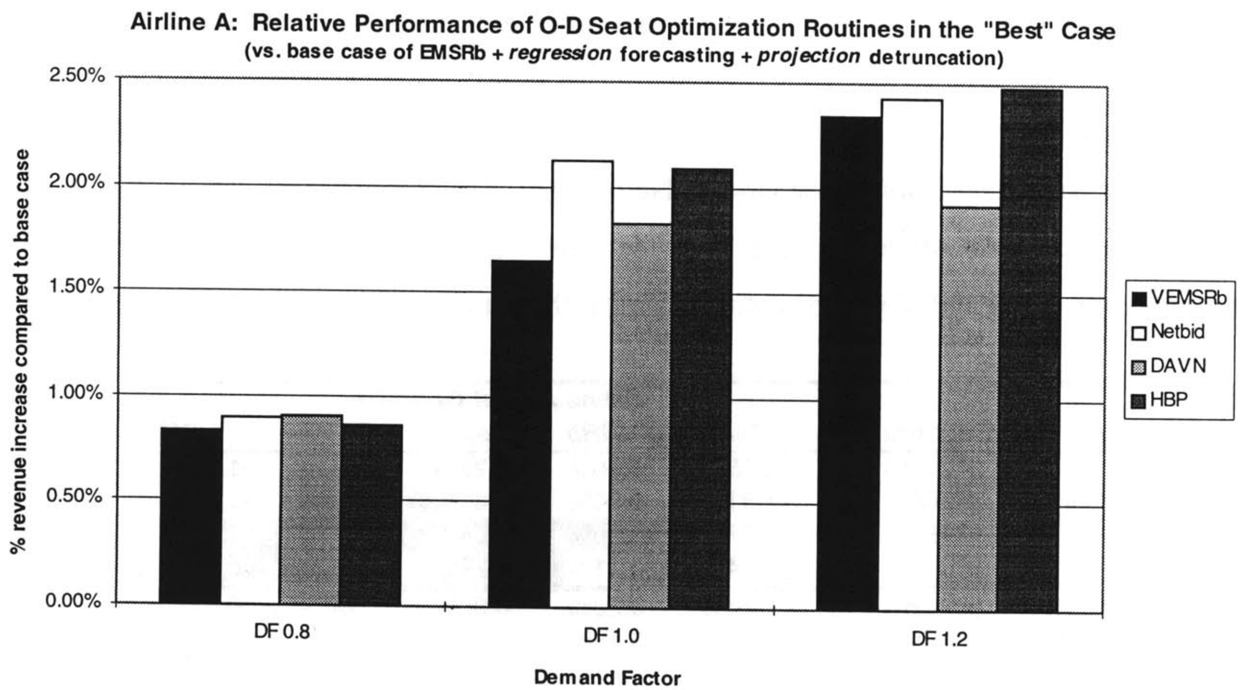

Figure 4.18: Revenue performance of O-D seat optimization routines in the "best" case

We can see that at the demand factor of 1.0, the gains exceed 1.5% in all cases of seat optimizer, and they are almost 2.5% (with the exception of DAVN) under a demand factor of 1.2. What this basically illustrates is that although our earlier results showed the revenue gains from the seat optimizers to be partially due to the less desirable forecasting and detruncation methods, even when the "best" case of forecaster and detruncator is implemented as the base case, significant revenue gain can still be achieved by implementation of O-D seat inventory control methods.

In terms of answering the question about the "best" set of algorithms, cases for which Airline B is using EMSRb have already been analyzed in Sections 4.1 and 4.2. Given that the competition (Airline B) is using EMSRb, Airline A's best choice has been shown to be Netbid when projection detruncation is used by Airline A, and DAVN when booking curve detruncation is chosen. These results reproduce themselves for all other options of forecasting and detruncation method when Airline B uses EMSRb seat optimization, and in fact, for all combinations of algorithm for either airline. Only the magnitude of the revenue gain/loss depends on the combination of the methodologies for the two airlines (refer to Appendix A.2 for the charts illustrating the revenue performances under all different combinations).

The best choices of forecasting and detruncation based on the analysis from Section 4.2.2 can also be presented for the airline of interest (Airline A) given its choice of seat optimization method. For Airline A using a given seat optimizer, its best combinations of forecaster and detruncation method are shown in Table 4.10 below.

Seat optimizer used	Forecast/detruncation best choice	Comments
EMSRb	regression/projection	---
VEMSRb	regression/projection	Any combination but *pickup/booking curve*
Netbid	regression/projection*	Do not use *booking curve* detruncation
DAVN	regression*/(proj. or b.c.)	*Pickup* ok vs. EMSRb, VEMSRb, Netbid
HBP	regression/projection	Any combination but *pickup/booking curve*

* indicates this appears to be the more important factor for revenue gains

Table 4.10: "Best" combinations of forecasting/detruncation methods

Another examination of the analysis can be taken in which the best choice for the seat optimizer can be determined given a selected set of forecasting/detruncation method inputs. Therefore, if Airline A will be using a given forecasting/detruncation combination, their optimal choice for a seat optimizer is given in Table 4.11 below.

Using forecast/detruncation combination	Best choice for seat optimizer	Comments
Pickup/booking curve	DAVN	---
Regression/booking curve	DAVN	HBP second, VEMSRb third
Regression/projection	Netbid (HBP ok)	DAVN/VEMSRb slightly lower but still very good
Pickup/projection	Netbid (HBP ok)	DAVN/VEMSRb slightly lower but still very good

Table 4.11: "Best" choices of seat optimization routines

In a general sense, it is evident that the combination of regression forecasting and projection detruncation consistently yields the highest revenue regardless of the seat optimization routine selected. However, if such a combination is not able to be implemented, the next best course of action can be taken based on the simulation results summarized in the preceding tables.

4.3 Changes in the τ parameter

In addition to an analysis of the passenger and revenue effects of forecasting/detruncation method combinations, another study was performed where the parameter of τ in the PODS input file was changed. Recalling from Section 2.2.2, τ is the user-defined guess for the percent of time by which we underestimate the demand given that a particular flight closed, and is illustrated and defined by Figure 4.19. Until now, an initial τ value of 0.15 has been used in all PODS 7b simulations; therefore, the purpose of this section is to analyze the results of changes in this parameter to ascertain the appropriateness of this value, and also to determine the sensitivity of the revenue results to the τ parameter.

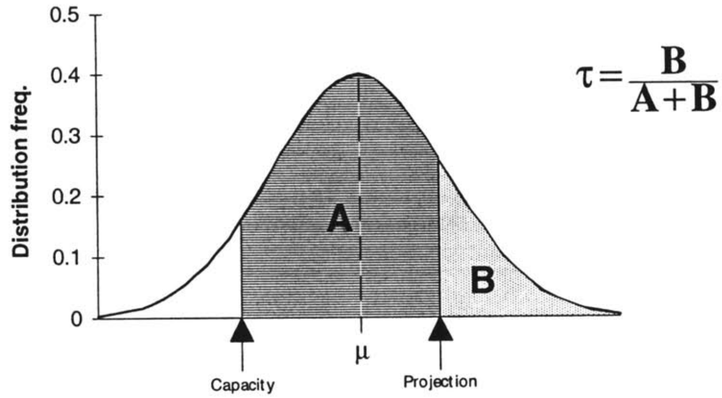

Figure 4.19: Constrained demand distribution for projection detruncation

Again, from Section 2.2.2 and Figure 4.19 above, τ is defined as the ratio of the shaded area B (i.e., demand is greater than what we projected) to the combined areas A+B (i.e., demand is greater than capacity on all closed observations). Hence, a τ value of 0.15 means that for the closed observations, our estimate of the projected demand was lower than the actual demand value 15% of the time. Establishing this mark along with the known capacity mark, the normal distribution permits us to thereby infer the projected demand, (labeled in Figure 4.19), for all given flights.

In simulations testing the sensitivity of projection detruncation to the τ parameter, τ was varied from an initial level of 0.15 to 0.25 and 0.35[76]. The obtained results indicate that changing τ from 0.15 to 0.25 or 0.35 affects system revenue performance differently under different seat optimizers, although the effect is limited. For example, with Airline B using the combination of EMSRb with pickup forecasting and booking curve detruncation (i.e., the full "base case" scenario), Airline A experienced a slight decrease in revenue both for $\tau = 0.25$ and for $\tau = 0.35$ when either EMSRb or VEMSRb was used. However, only minimal changes were encountered using any of the other seat optimization methods (i.e., Netbid, DAVN, and HBP). These results can be seen in Figure 4.20, where the combinations of regression/projection and pickup/projection were the tested methods of forecasting/detruncation for Airline A. The maximum change in

[76] $\tau = 0.15$ had been used in all previous simulations.

revenues due to variation of the parameter τ from 0.15 to 0.35 (i.e., the maximum variation between any set of grouped data columns) was 0.6% or less, with the exception of EMSRb, which had variation on the order of 1.3-1.4%.

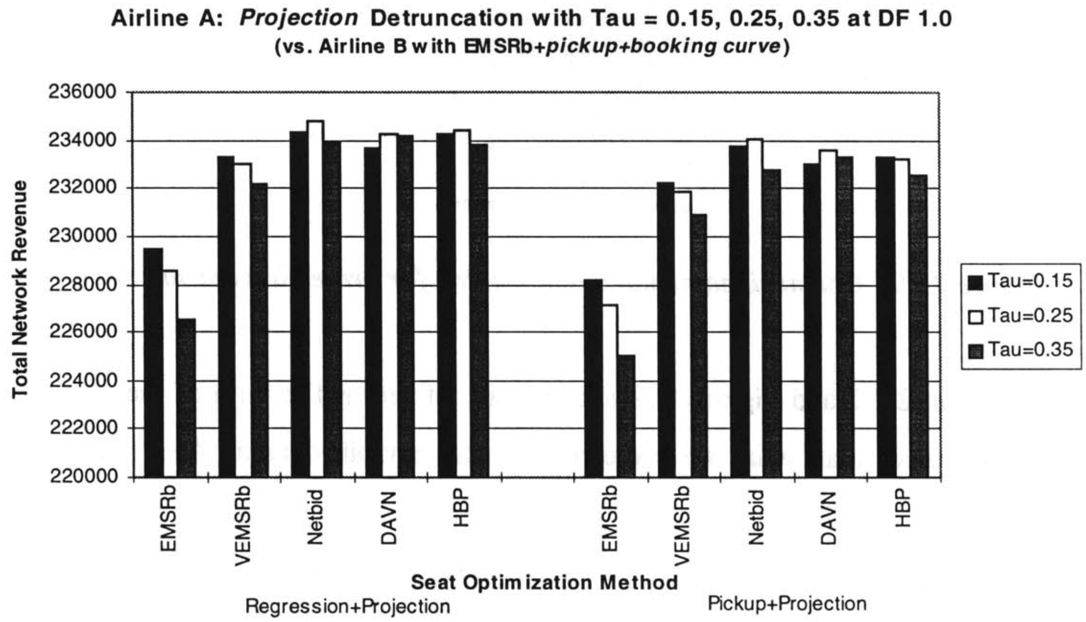

Figure 4.20: Effect of τ with Airline B using booking curve detruncation

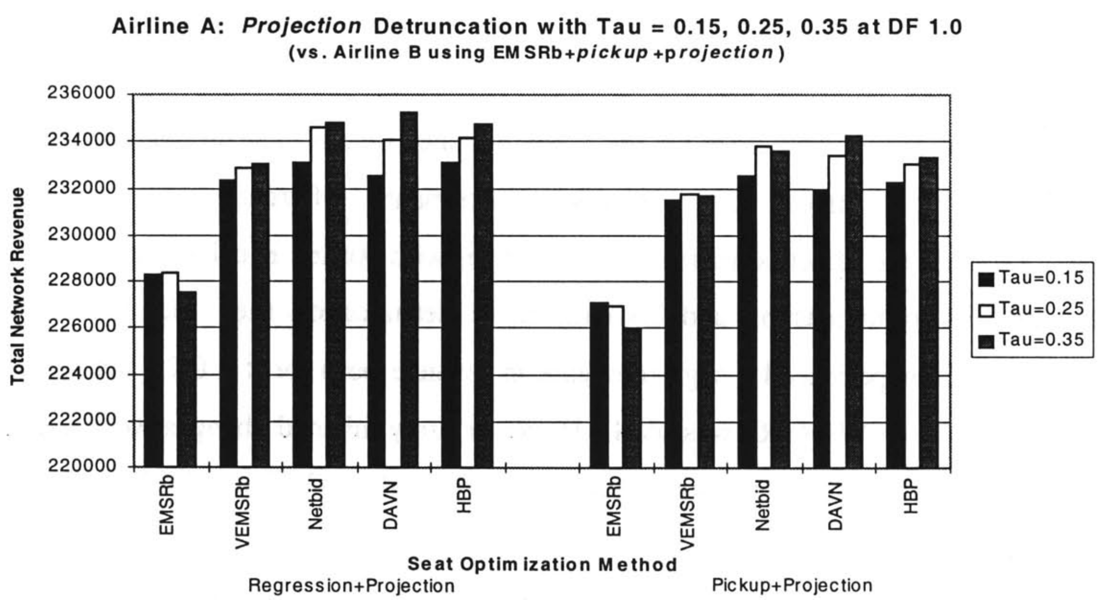

Figure 4.21: Effect of τ with Airline B using projection detruncation

The aforementioned results compared the effect of changes in τ for Airline A only, as Airline B was using booking curve detruncation. If we therefore change Airline B's methodology to EMSRb with pickup forecasting and *projection detruncation*--so that both airlines' projection detruncation methods are now sensitive to changes in τ--a new set of results is obtained, shown in Figure 4.21.

Here we see that EMSRb and VEMSRb are not as sensitive to changes in τ, while Netbid, DAVN, and HBP appear to be slightly more impacted by them (Netbid appears to have the largest sensitivity to changes in τ). This most likely explanation for this is because Airline B is now using projection detruncation, which causes their revenue to increase compared to what was experienced under booking curve detruncation; some of this improvement comes at the expense of Airline A's revenue (note that the absolute revenue values are lower in Figure 4.21 than in Figure 4.20). Hence, in the former case, Airline B was not using projection detruncation so Airline A was capturing higher revenues, regardless of the choice of the τ value (i.e., the τ value was not as important because Airline A had a superior detruncation method to begin with). However, in the latter case of a more competitive situation (i.e., both airlines using projection detruncation), such changes in τ will now have larger impacts on network revenues. This is corroborated by the fact that the seat optimization routines which perform well under projection detruncation (i.e., Netbid, DAVN, and HBP) are those in which a slight revenue increase is seen for increasing τ in the latter case (Figure 4.21). In any case, the magnitude of the variation induced by changes in the τ parameter from 0.15 to 0.35 is still on the order of 0.5%, with the exception of DAVN, for which it barely exceeds 1.0%.

4.4 Factors contributing to revenue increases

The final section of this chapter incorporates the trends observed in the preceding sections to provide some explanation for the factors behind the increases in revenues observed. More specifically, we have seen that regression forecasting generally results in higher revenues, as does implementation of projection detruncation. These revenue increases have resulted from a better passenger mix due to better fare class protection and higher forecasted demand. However, the reasons why the regression forecasting and projection detruncation algorithms cause higher forecasts of demand have not been analyzed. Therefore, the purpose of the following subsections is to lend insight into the performance of the different forecasting, detruncation, and seat optimization method themselves.

4.4.1 Explanation of forecasting/detruncation method performance

Although both regression forecasting and projection detruncation have resulted in overall system revenue increases, the gains from a change to projection detruncation have been much larger than those from a change to regression forecasting. This is explained by comparing the forecasts of the different methods. In essence, the regression forecaster is more similar to its pickup counterpart in forecasted demands than the projection detruncator is to its booking curve counterpart (see Sections 4.1.2.2 and 4.1.2.4 for discussions of passenger loads and forecasted demands, respectively).

Another question can be asked as to why regression forecasting outperforms pickup forecasting. Regression forecasting should theoretically perform better than simple pickup forecasting, as pickup forecasting is just a constrained form of regression forecasting, as explained in Section 2.1.2. Allowing an extra degree of freedom by allowing the slope coefficient to vary should help make a more accurate forecast.

A third question can be asked comparing the detruncation methods--why does projection detruncation outperform booking curve detruncation? Projection detruncation was seen to perform better than booking curve detruncation with each of the seat optimization routines, primarily due to the fact that higher forecasted demands were realized. But why is this the case? It is inherent in the way in which booking curve detruncation unconstrains the demand--for the periods in which a fare class is closed, the trends on unclosed flight observations are applied to this period of the closed flights. However, the simple fact that a flight was unclosed means that demand was lower than projected, and therefore, the booking curve multiplier ratios are only truly reflective of these lower-demand flights and are probably too low to be applied to flights with high demand (i.e., those that closed at some point in the booking process). Application of these ratios to the higher-demand flights generally results in extrapolations of the booking data which produce forecasts that are too low. Projection detruncation addresses this problem by using a demand distribution and projecting demand on closed flight observations based on an assumption of how often we underpredict the actual demand. This is done by estimating an average demand which incorporates *both* closed and unclosed flights until this mean value exhibits convergence.

Furthermore, another complicating factor of the actual accuracy of the detruncation routines is the number of times which detruncation is actually performed. As demand becomes higher, fare classes tend to close more often, and therefore more detruncation needs to be done. So a method which tends to underpredict demand as it unconstrains it (as booking curve detruncation does) will thereby need to perform detruncation more frequently as the demand becomes higher.

4.4.2 Explanation of seat optimization algorithm performance

As for the seat optimization algorithms themselves, we have seen that most have performed in the network case according to our initial hypotheses, with the exception of Netbid. The EMSRb algorithm should theoretically have the lowest performance, since it is allocating seats on a leg basis simply by fare class designation, with no adjustment for fare class value or passenger displacement costs. VEMSRb is an improvement on this and should therefore perform better as it is accounting for fare class value in determining which itineraries to accept/reject, but since no displacement costs are being calculated, it is still "greedy" and not aptly suited for implementation on an O-D network.

DAVN and HBP, on the other hand, were the methods which resulted in the most consistent overall performance. HBP performs well in that it accounts for passenger displacement costs using a bid price scheme with EMSR technology. It compares the EMSRs of each leg of the itinerary while also accounting for the fact that a local customer may be displaced on either of these legs. DAVN does well because it is a method which ranks itineraries by fare class value, while also accounting for passenger displacement costs by calculating pseudo fares derived from shadow prices (something that VEMSRb fails to do). Hence, it too is a "greedy" scheme (but not overly "greedy"), as passenger displacement costs are accounted for.

Upon the initial results of Netbid's poor performance under the base case forecaster and detruncator, possible reasons for its poor performance were put forth; namely, that it was unable to perform well on such a small network, that solving the deterministic linear program was not optimal, and that the forecasting and detruncation methods being used were not compatible with the seat optimizer itself. This third reason has been examined within the scope of the thesis, and it has been found that such reasoning was correct--Netbid performs poorly when the forecasted demands are too low. Why is this? The answer lies in the way in which Netbid controls seat inventory--it does this using a bid price scheme, with which the control is not as tight (refer to Section 4.1.2.4 for an example). Hence, while the current simulation still has Netbid solving a deterministic linear program on a small network, revenue gains commensurate with (and exceeding) those achieved with other O-D seat inventory control methods were able to be obtained with a better choice of forecaster and detruncator.

A final reason for the better revenue performance of the forecasting and detruncation methods is that of competitive feedback effects within the PODS simulator. While most traditional simulators model the airline booking process with independent demands by fare class and for each carrier, the PODS model has full passenger choice of path and carrier, as explained in Section 3.1. Competitive feedback effects therefore play a large role in accentuating the revenue gains or losses under a variety of yield management methods. For instance, if Airline A's combination of forecaster, detruncator, and seat optimization algorithm is superior to that of Airline B, we would expect the revenues for Airline A to increase more drastically due the interaction between the two competitors. That is, if the yield management system for Airline A can do a better job of rejecting low-fare Q-class demand and protecting seats for the high-fare Y-class demand (assuming this demand materializes), it will tend to spill a larger number of Q-class passengers. Given that Airline B has a yield management system which is less able to reject such customers, it will in turn fail to reject many of these requests and thereby fill more of its aircraft with a lower-revenue passenger mix. And since we know that the lower-fare customers generally tend to book reservations earliest in the process, acceptance of too many of these customers by Airline B will subsequently cause spill of higher fare customers later in the booking process--customers which are then spilled to Airline A and become extra Y-class demand (for which seats will have been protected). High acceptance of such passengers may then cause Airline A's Y-class to close down more often, for which the detruncation algorithm will further increase its subsequent demand forecasts, and this cycle will then be repeated.

4.4.3 Comparison of results with prior findings

Up to this point, few studies have examined the relative impact of O-D seat inventory control methods under different forecasting and detruncation methods, so direct comparisons cannot be made. However, a selected number of papers have been written comparing the performance of forecasting (and detruncation) methods, of which three have results from which interesting comparisons to this research can be drawn.

First, Williamson[77] examined different network seat inventory control models (some of which were tested here), and tested their performance relative to simple EMSRb seat optimization. The primary item of note is the performance of Netbid--it resulted in relatively little revenue gain over the EMSR algorithm, except at very high load factors. Second, Wickham[78] compared several different forecasting models and measured their accuracy, using actual airline data. Notable results were that pickup forecasting generally outperformed regression except in the highest booking classes (this is contrary to what was found in this thesis). However, the simulator used by Wickham had no competitive feedback effects present, which could change the relative performance of these methods. Furthermore, a general positive bias was found in almost all forecasting methods, meaning that the forecasted values were higher than the actual demands. Since forecast accuracy is not able to be measured under a correlated demand situation as in PODS, no similar comparison with actual demands can be done. However, it was indeed seen that there was a positive bias of the regression forecaster and the projection detruncator relative to the base case methodologies of pickup forecasting and booking curve detruncation, respectively (i.e., the regression and projection methods always produced higher forecasts).

---

[77] Williamson (1992), see Section 6.1.3.

[78] Wickham (1995). Section 4.4 and Chapter 5 provide general summaries of the results.

Finally, Skwarek[79] tested forecasting and detruncation impacts on the PODS simulator itself, although these tests were done for the single-leg case with only EMSRb seat inventory control technology. For the forecasting methods tested, he found that regression forecasting outperformed pickup, except in the highest-demand cases (directly the opposite effect seen by Wickham!); the lower-demand cases match with what was found in this thesis. In terms of detruncation results, Skwarek found that projection and booking curve performed about the same, except at the highest demand factors, in which the higher unconstraining demand on closed flights caused much higher revenue for projection detruncation. This phenomenon was accentuated in the PODS results in this thesis when both detruncation methods were tested in combination with the O-D seat optimizers.

Basically, this chapter has presented the general results of the PODS simulations tested and analyzed in regards to forecasting and detruncation impacts in combination with the seat optimizers. And while we have seen that while previous results and conclusions about the seat optimizers are still valid, these results may change due to different forecasting/detruncation methods, given the volatility of the different yield management algorithms to such inputs. The observed trends remained more or less the same regardless of the competitive response of the competitor, and best-case scenarios of methodologies for different competitive situations were presented. A general summary of the findings is the primary focus of the next chapter.

[79] Skwarek (1996), see Sections 7.2.1 and 7.2.2.

Chapter 5

Conclusions

The fifth and final chapter of this thesis provides a review of its objectives along with a summary of the results obtained in the PODS simulations performed. Finally, questions still unanswered by this research are posed, from which possible or interesting directions for new research emerge.

5.1 Review of thesis objectives

The preceding chapters of this thesis have focused on a variety of different methods for forecasting and detruncation as applied to airline revenue management using the PODS simulator developed at Boeing. The first chapter provided a brief introduction to revenue management and its forecasting and detruncation components, while the second chapter discussed in detail the different forecasting and detruncation methods tested, along with shorter descriptions of the seat optimization routines tested in this latest version of PODS. The third chapter then provided a description of the PODS simulator and its inputs, from overall market-level inputs (e.g., the route network structure) to airline-specific inputs (e.g., the choice of forecasting methodology); and output performance metrics were also discussed. The fourth chapter then explored the impacts resulting from variations in combinations of the PODS inputs; primarily those of forecasting, detruncation, and seat optimization methodology under three different demand factors.

As discussed, there were two basic objectives for the experiments undertaken. The first was to expand previous PODS-related research to a network scenario (here, six individual spoke cities, two connecting airport hubs, and two competing airlines) and analyze the effects of different forecasting and detruncation algorithms in the yield management context. These results could then be compared with earlier revenue management simulations in order to determine the sources of gains from yield management system improvement (i.e., from the seat optimization algorithm itself or rather the choice of forecaster/detruncator). The second was to simulate a more real-world network scenario, where each competitor has the ability to vary its choice of seat optimization, forecasting, and detruncation method, and to thereby determine which forecasting/detruncation methods perform best in combination with the different seat optimization routines tested.

From these objectives came two primary goals: (1) to obtain a set of simulated airline network results highlighting the differences among combinations of forecaster, detruncator, and seat optimizer while gaining insight into why different seat optimization systems perform differently under the different methods, and (2) to find a "best-case" combination of algorithms to determine the best network revenue management options under a variety of different competitive scenarios. These goals were achieved based upon various analyses of different output performance measures; namely, system revenues, passenger loads, fare class closure rates, and forecasted demands.

5.2 Summary of findings and general results

The objectives described above were achieved by first analyzing scenarios in which the airline of interest had the ability to vary its forecasting and detruncation algorithms while the competing airline was held to a base case set of algorithms, and later moving to scenarios where both airlines were free to vary their forecasting, detruncation, and seat optimization methods.

Initially, system revenue values for Airline A were analyzed to gain insight into the relative overall performance of the various combinations of methods. Previously obtained PODS network results with a base case of pickup forecasting and booking curve detruncation at a demand factor of 1.0 were reproduced[80], and subsequent comparisons made between this base case and other yield management scenarios with different combinations of algorithms. Such results showed that while two of the four O-D seat inventory control methods tested (i.e., DAVN and HBP) performed rather well by having system revenue gains on the order of 2.0-2.5% over the base case, the Netbid seat optimizer ended up with rather disappointing revenue percentage gains at less than 1.0% over the same base case.

[80] Lee (1998). Refer to Section 4.1.3 for results.

Changes were then implemented in the forecasting and detruncation methods for the airline of interest, to see whether the relative ranking of O-D methods remained consistent at a demand factor of 1.0 (recall that Airline B was still using EMSRb with pickup forecasting and booking curve detruncation). While the DAVN and HBP revenue gains over the base case remained relatively the same (still in the range of 2.5-3.0%), indicating little sensitivity to the choice of forecasting or detruncation methodology; the gains for Netbid were seen to increase drastically in the cases where projection detruncation was implemented (to more than 3.0% over the base case). Hence, Netbid became the method with the highest simulated revenue, given the same set of competitive conditions when projection detruncation was used, while it had been performing quite poorly when booking curve detruncation was used.

One could ask whether these gains were being inflated by the fact that the base case was always using the least revenue-beneficial EMSRb seat optimization. Therefore, further analysis compared the results to a base case of pickup and booking curve using the same seat optimizer at a demand factor of 1.0, under which it was seen that a change to regression forecasting from pickup increased revenues by about 0.5%, regardless of the seat optimization method. However, the change from booking curve to projection detruncation results in increases of approximately 0.5-1.0% except for Netbid, which saw gains as high as 2.0%.

From these analyses, conclusions were drawn about the "robustness" of the various methods tested. It was seen that DAVN is a robust method, and HBP performs rather consistently also. However, Netbid is quite susceptible to the choice of detruncation method; poor performance resulted from the use of booking curve detruncation while excellent performance occurred when projection detruncation was used. The overall best choice for forecasting and detruncation was regression/projection, regardless of the seat optimization method, while the best choice for seat optimization was Netbid if projection detruncation was used and DAVN when booking curve detruncation was chosen.

Flight leg passenger loads broken down by fare class were also analyzed to see what changes took place under the different forecasting and detruncation methods, and to provide insight into the revenue results summarized above. In the base case scenario of pickup forecasting and booking curve detruncation, it was seen that Netbid's leg load factors were appreciably higher than those under the other seat optimization routines, primarily due to extra Q-class (low-fare) passengers. However, once the forecasting and detruncation methods were changed, passenger load changes also occurred. In changing the forecaster from pickup to regression, the leg load changes were minimal under both booking curve and projection detruncation. But when the detruncator was changed from booking curve to projection (under either pickup or regression forecasting), the changes in the passenger leg loads were more pronounced. The Y-class (high-fare) loads increased, while the Q-class (low-fare) loads decreased at 2-3 times the magnitude, resulting in overall lower total leg loads (the B- and M-class loads were relatively small and changes in them were negligible). These effects were most pronounced for Netbid. Hence, the higher revenues seen were generally a result of a better passenger mix due to increased rejection of lower fare passengers and more seat protection for higher-fare customers, as will be discussed next.

Because the driving force of passenger loads is fare class closures, these were subsequently examined. It was seen that the regression/projection combination caused earlier fare class closures than the pickup/booking curve combination for all cases of seat optimization. In fact, with the regression/projection combination, only Y-class fares became available for a large portion of the booking process on a sampled high demand leg, while the lowest classes (M, Q) were almost always closed under best method combinations (at least in the higher demand scenarios). Therefore, this forecasting/detruncation combination is able to reject lower fare classes better than the pickup/booking curve combination.

To wrap up the base case analysis, the forecasted remaining demands were examined for the cases of pickup/booking curve and regression/projection. While difficult to interpret for seat optimization methods using virtual classes (as different ODFs map to different virtual buckets), EMSRb and Netbid showed pronounced differences between the two combinations. With regression/projection, forecasts were higher in all fare classes, although those in Y- and Q-class were much larger than those in B- and M-class. It should be noted that the forecasts in the highest fare class are the most important and have the largest influence on revenues, as top-down seat protection is used and seats protected accordingly. So larger Y-class forecasts will produce higher revenues if the extra protected seats can be filled (i.e., if there is enough demand). In the simulations here, it is seen that pickup/booking curve produced forecasts which were too low. When these forecasts were increased by using regression/projection, the demand was sufficient to fill the extra protected seats, resulting in a better passenger mix and higher system revenues.

All of the above results were in scenarios where Airline B was fixed using its base case of EMSRb with pickup forecasting and booking curve detruncation. Additional tests where the two competitors have the same seat optimizer were done, although the trends appeared similar; that is, most of the revenue gain still comes from a better forecasting/detruncation combination, with Netbid having the most pronounced increase of the different seat optimizers. When full flexibility and choice of forecasting, detruncation, and seat optimization algorithm was given to both competitors, the relative rankings still held. Regression/projection was always the best combination, regardless of seat optimization choice or competitive situation; otherwise, the choice of seat optimizer should be made according to the detruncation method used. Netbid performs very well with projection detruncation, although DAVN and HBP do only slightly worse, while DAVN and HBP are the best choices if booking curve detruncation is to be used. Lastly, reasons behind the performance of the different methods themselves were put forth.

5.3 Unanswered questions

Although this thesis has presented a rather comprehensive set of network simulation results from PODS which relate forecasting and detruncation impacts to airline network revenue management, there are still a number of questions of interest yet unanswered by the analysis of results presented in Chapter 4. Some of these are discussed below.

First, what were the actual passenger demands by flight leg or market? In practice, this is impossible to determine, given that we are modeling correlated demands with passenger path choice, as the demand by ODF is itself a function of the path availability. However, knowing the actual demands would allow for comparison with the forecasted demands from which determinations could then be made as to which forecasting method was most accurate under the different seat optimizers and why. That is, are the revenues highest for a given forecasting method because those forecasts were the most accurate, or rather because they were favorably overpredicted or underpredicted (i.e., maybe it is better to consistently overpredict/underpredict demand under certain demand conditions)?

Second, what were the actual values determined by each individual forecasting and detruncation algorithm and how did these values compare across methods? Knowing this would allow a relative comparison of the methods themselves and would lend some insight into the results building upon the first question above; namely, the reasoning behind why revenues were higher. The relatively good performance of projection detruncation could have resulted from too little or too much detruncation, or instead because booking curve detruncation produced results that were just too low.

Third, would the results change in other network cases; for example, larger networks? The PODS simulation here was performed for a network with six spoke cities, although a six-city network could still produce strikingly different results than one that encompasses hundreds of cities, as seen in the domestic US industry. Of interest, therefore, would be to see whether yield management methods perform differently when size of network is a consideration. Furthermore, what effects result when the demands are spread randomly over the network, as opposed to the judgmental manner in which demand was set for this simulation[81]?

Finally, what effects appear from the use of other yield management methods which were not currently available in PODS, and how do these perform? Two examples are presented by Wei--the non-greedy heuristic bid price model and the convergent EMSR model[82]. The former can be thought of as a combination of DAVN and HBP, in which pseudo fares are calculated as in DAVN and then used to solve for the EMSR values, after which the bid prices are calculated using these EMSR heuristics as in HBP, while the latter tries to come up with a better allocation

[81] Demand could also be assigned randomly to legs throughout the network rather than setting the short-haul legs to have the highest *a priori* demand.
[82] Wei (1997). Sections 3.1 and 3.2 describe each of these two methods.

of the passenger's fare across the legs they traverse before making the EMSR calculations as in the EMSRb method[83]. The effects of implementation of these methods or other ones not tested here and their interaction with the different forecasting and detruncation routines would also be of interest.

5.4 New research directions

The questions left unanswered by this research, discussed above in Section 5.3, lead the reader to ponder new directions for future research in this area. Hence, several new research questions which might be of interest are discussed below.

First, what is the actual effect of larger networks, as discussed above? Are the PODS results similar in these cases or do they differ? Preliminary tests[84] have been performed in PODS on a ten-city network, although no complete testing of the forecasting and detruncation routines under such a scenario has been done as of yet. Second, the possibility of using other forecasting/detruncation methods in PODS is of interest, especially that of the efficient forecaster[85]. Questions of interest would be those detailing how it compares with the methods tested in this thesis. A third direction for future research could also be that of testing other seat optimization methods (e.g., the convergent EMSR model), as discussed in Section 5.3. Again of interest would be their relative performance in combination with the forecasting/detruncation methods used. Fourth, changes or improvements in the forecasting, detruncation, and seat optimization methods themselves could be studied. One example occurs in projection detruncation, where the parameter τ is manually chosen and is constant for all paths. However, the methodology could be modified to have τ be determined on a path basis as a function of the booking curve and the time frame in which the path fare class closed. Fifth, modified networks which allow for flights between the hub airports or have more than two competitors/hubs are also of interest. The former scenario gives multiple path alternatives to what was classified as local demand (i.e., those passengers traveling from hub to spoke or vice versa who previously had only

[83] Also see Bratu (1998) for more detailed discussion on the convergent EMSR method applied at an O-D level.
[84] Lee (1998). Similar relative trends as presented in this thesis have been seen in initial cases run under the ten-city network, although Netbid's poor performance is not as pronounced as in the six-city case.
[85] Refer to Skwarek (1996), Section 5.2.3 for a description of the efficient forecaster.

one path choice), while the latter would show whether revenue gains were of the same magnitude as the number of path choices increases for each ODF combination. Sixth, more than one connecting bank can be used. In the current PODS formulation, only a single flight departure was used in each market, whereas multiple connecting banks allow a time-dependent component of path choice to occur. Finally, booking cancellations and no-shows can be permitted and the overbooking concept can be implemented, to more closely model a real-world network case. All of these can be studied not only in terms of the seat optimization routines, but also in terms of their interaction with the forecasting and detruncation algorithms presented.

Such research questions as those above can be tested under the current PODS methodology, given some implementation changes and input variations. However, a more broad direction for future research would be one that moves a step closer toward complete optimization of the airline supply problem (i.e., provide interactive optimization of schedules, yield management, and fares). However, with current technology this proves extremely difficult, due to several considerations. First, the sheer size of the problem and the associated billions of decision variables and equations would render the problem too large to be solved in a reasonable amount of time on a repeated basis. Even if this were possible, the competitive response is very difficult to predict, hence, any "optimal" scenario would be assuming a set of fixed competitive characteristics (or responses). Furthermore, the many constraints such as fleet considerations, airport restrictions, etc. only make the task more complex. Despite these difficulties, two research directions which could be brought to fruition in the near future are those of flexible pricing and ODF grouping. Flexible pricing means the pricing regime used by an airline is incorporated and solved along with the seat allocation problem. Instead of having a fixed pricing structure, fares are optimized simultaneously with seat allocations given knowledge of future demands, rendering the pricing structure itself dependent on the expected ODF demand. ODF grouping[86] can be used to combat the "small numbers" problem so that ODF forecasts used during seat optimization are more reliable and more accurate. An example of this would be to group distinct O-D markets which possess similar characteristics, while another case would entail aggregation of different flights by time (i.e., one se of booking limits could be established for two or more consecutive flights). Such research

[86] Williamson (1992). Section 4.3 discusses aggregation of ODFs.

directions would provide the next improvements in ODF forecasting, which would hopefully continue to be beneficial to revenues.

In any case, large advances have been made over the past couple of decades in the airline supply problem, with implementations such as yield management (first leg and now O-D based), fleet assignment, and crew scheduling saving millions of dollars for many major airlines worldwide. Continued research on a model such as PODS can only assist in further removing inefficiencies by which airlines are constrained today.

Appendix A

A.1: PODS airline revenue results vs. competitor with pickup forecasting, booking curve detruncation, and EMSRb seat optimization, at demand factors 0.8, 1.0, and 1.2

A.2: PODS airline revenue results vs. all competitor possibilities at demand factor 1.0

A.1: PODS airline revenue results vs. competitor with pickup forecasting, booking curve detruncation, and EMSRb seat optimization, at demand factors 0.8, 1.0, and 1.2

	------Airline A------			------Airline B------			*Revenues*		*% increase compared to EMSRb*	
DF	**Seat Optimizer**	**Forecast**	**Detruncation**	**Seat Optimizer**	**Forecast**	**Detruncation**	**Airline A**	**Airline B**	**Airline A**	**Airline B**
0.8	EMSRb	Pickup	Book. Curve	EMSRb	Pickup	Book. Curve	189,813	189,703	-----	-----
0.8	VEMSRb	Pickup	Book. Curve	EMSRb	Pickup	Book. Curve	191,115	188,610	0.69%	-0.58%
0.8	Netbid	Pickup	Book. Curve	EMSRb	Pickup	Book. Curve	190,076	189,174	0.14%	-0.28%
0.8	DAVN	Pickup	Book. Curve	EMSRb	Pickup	Book. Curve	191,365	188,497	0.82%	-0.64%
0.8	HBP	Pickup	Book. Curve	EMSRb	Pickup	Book. Curve	191,016	188,698	0.63%	-0.53%
0.8	EMSRb	Reg.	Book. Curve	EMSRb	Pickup	Book. Curve	190,482	189,556	0.35%	-0.08%
0.8	VEMSRb	Reg.	Book. Curve	EMSRb	Pickup	Book. Curve	191,733	188,405	1.01%	-0.68%
0.8	Netbid	Reg.	Book. Curve	EMSRb	Pickup	Book. Curve	191,013	188,853	0.63%	-0.45%
0.8	DAVN	Reg.	Book. Curve	EMSRb	Pickup	Book. Curve	191,938	188,317	1.12%	-0.73%
0.8	HBP	Reg.	Book. Curve	EMSRb	Pickup	Book. Curve	191,707	188,503	1.00%	-0.63%
0.8	EMSRb	Reg	Proj.	EMSRb	Pickup	Book. Curve	190,703	189,386	0.47%	-0.17%
0.8	VEMSRb	Reg.	Proj.	EMSRb	Pickup	Book. Curve	192,295	187,972	1.31%	-0.91%
0.8	Netbid	Reg.	Proj.	EMSRb	Pickup	Book. Curve	192,406	188,067	1.37%	-0.86%
0.8	DAVN	Reg.	Proj.	EMSRb	Pickup	Book. Curve	192,423	187,907	1.38%	-0.95%
0.8	HBP	Reg.	Proj	EMSRb	Pickup	Book. Curve	192,342	188,083	1.33%	-0.85%
0.8	EMSRb	Pickup	Proj.	EMSRb	Pickup	Book. Curve	190,098	189,525	0.15%	-0.09%
0.8	VEMSRb	Pickup	Proj.	EMSRb	Pickup	Book. Curve	191,789	188,133	1.04%	-0.83%
0.8	Netbid	Pickup	Proj.	EMSRb	Pickup	Book. Curve	191,939	188,145	1.12%	-0.82%
0.8	DAVN	Pickup	Proj.	EMSRb	Pickup	Book. Curve	192,065	187,973	1.19%	-0.91%
0.8	HBP	Pickup	Proj	EMSRb	Pickup	Book. Curve	191,811	188,206	1.05%	-0.79%
1.0	*EMSRb*	*Pickup*	*Book Curve*	*EMSRb*	*Pickup*	*Book Curve*	*226,954*	*226,952*	-----	-----
1.0	VEMSRb	Pickup	Book. Curve	EMSRb	Pickup	Book. Curve	230,529	224,804	1.58%	-0.95%
1.0	Netbid	Pickup	Book. Curve	EMSRb	Pickup	Book. Curve	228,871	226,112	0.84%	-0.37%
1.0	DAVN	Pickup	Book Curve	EMSRb	Pickup	Book. Curve	232,231	224,225	2.33%	-1.20%
1.0	HBP	Pickup	Book. Curve	EMSRb	Pickup	Book. Curve	231,414	224,686	1.97%	-1.00%
1.0	EMSRb	Reg.	Book. Curve	EMSRb	Pickup	Book. Curve	227,859	227,191	0.40%	0.11%
1.0	VEMSRb	Reg.	Book. Curve	EMSRb	Pickup	Book. Curve	231,647	224,580	2.07%	-1.05%
1.0	Netbid	Reg.	Book. Curve	EMSRb	Pickup	Book. Curve	229,955	225,959	1.32%	-0.44%
1.0	DAVN	Reg.	Book. Curve	EMSRb	Pickup	Book Curve	233,311	224,048	2.80%	-1.28%
1.0	HBP	Reg.	Book. Curve	EMSRb	Pickup	Book Curve	232,400	224,561	2.40%	-1.05%
1.0	EMSRb	Reg	Proj.	EMSRb	Pickup	Book Curve	229,494	225,959	1.12%	-0.44%
1.0	VEMSRb	Reg	Proj.	EMSRb	Pickup	Book Curve	233,276	223,588	2.79%	-1.48%
1.0	Netbid	Reg.	Proj.	EMSRb	Pickup	Book Curve	234,378	223,602	3.27%	-1.48%
1.0	DAVN	Reg.	Proj.	EMSRb	Pickup	Book. Curve	233,687	223,591	2.97%	-1.48%
1.0	HBP	Reg.	Proj.	EMSRb	Pickup	Book Curve	234,287	223,251	3.23%	-1.63%
1.0	EMSRb	Pickup	Proj	EMSRb	Pickup	Book. Curve	228,210	226,042	0.55%	-0.40%
1.0	VEMSRb	Pickup	Proj	EMSRb	Pickup	Book. Curve	232,250	223,783	2.33%	-1.40%
1.0	Netbid	Pickup	Proj	EMSRb	Pickup	Book. Curve	233,744	223,591	2.99%	-1.48%
1.0	DAVN	Pickup	Proj.	EMSRb	Pickup	Book. Curve	233,023	223,602	2.67%	-1.48%
1.0	HBP	Pickup	Proj.	EMSRb	Pickup	Book Curve	233,279	223,358	2.79%	-1.58%
1.2	*EMSRb*	*Pickup*	*Book Curve*	*EMSRb*	*Pickup*	*Book Curve*	*259,762*	*260,029*	-----	-----
1.2	VEMSRb	Pickup	Book. Curve	EMSRb	Pickup	Book Curve	263,265	258,375	1.35%	-0.64%
1.2	Netbid	Pickup	Book. Curve	EMSRb	Pickup	Book Curve	262,526	260,088	1.06%	0.02%
1.2	DAVN	Pickup	Book. Curve	EMSRb	Pickup	Book. Curve	266,801	257,353	2.71%	-1.03%
1.2	HBP	Pickup	Book Curve	EMSRb	Pickup	Book. Curve	266,292	257,339	2.51%	-1.03%
1.2	EMSRb	Reg	Book. Curve	EMSRb	Pickup	Book Curve	260,626	260,915	0.33%	0.34%
1.2	VEMSRb	Reg	Book. Curve	EMSRb	Pickup	Book. Curve	263,980	258,449	1.62%	-0.61%
1.2	Netbid	Reg	Book. Curve	EMSRb	Pickup	Book. Curve	262,950	260,520	1.23%	0.19%
1.2	DAVN	Reg.	Book Curve	EMSRb	Pickup	Book. Curve	268,247	257,149	3.27%	-1.11%
1.2	HBP	Reg.	Book. Curve	EMSRb	Pickup	Book. Curve	267,634	257,330	3.03%	-1.04%
1.2	EMSRb	Reg	Proj	EMSRb	Pickup	Book. Curve	263,348	258,928	1.38%	-0.42%
1.2	VEMSRb	Reg	Proj	EMSRb	Pickup	Book. Curve	269,526	255,793	3.76%	-1.63%
1.2	Netbid	Reg	Proj.	EMSRb	Pickup	Book. Curve	269,760	256,571	3.85%	-1.33%
1.2	DAVN	Reg	Proj.	EMSRb	Pickup	Book. Curve	268,407	257,300	3.33%	-1.05%
1.2	HBP	Reg	Proj.	EMSRb	Pickup	Book Curve	269,887	255,857	3.90%	-1.60%
1.2	EMSRb	Pickup	Proj	EMSRb	Pickup	Book. Curve	261,299	258,877	0.59%	-0.44%
1.2	VEMSRb	Pickup	Proj	EMSRb	Pickup	Book Curve	267,946	256,050	3.15%	-1.53%
1.2	Netbid	Pickup	Proj.	EMSRb	Pickup	Book. Curve	268,631	256,544	3.41%	-1.34%
1.2	DAVN	Pickup	Proj	EMSRb	Pickup	Book. Curve	267,344	257,219	2.92%	-1.08%
1.2	HBP	Pickup	Proj	EMSRb	Pickup	Book. Curve	268,345	255,926	3.30%	-1.58%

A.2: PODS airline revenue results vs. all competitor possibilities at demand factor 1.0

	-----Airline A-----			-----Airline B-----			Revenues		% increase compared to EMSRb	
DF	Seat Optimizer	Forecast	Detruncation	Seat Optimizer	Forecast	Detruncation	Airline A	Airline B	Airline A	Airline B
1.0	*EMSRb*	*Pickup*	*Book. Curve*	*EMSRb*	*Pickup*	*Book. Curve*	*226,954*	*226,952*	----	----
1.0	VEMSRb	Pickup	Book. Curve	EMSRb	Pickup	Book. Curve	230,529	224,804	*1.58%*	*-0.95%*
1.0	Netbid	Pickup	Book. Curve	EMSRb	Pickup	Book. Curve	228,871	226,112	*0.84%*	*-0.37%*
1.0	DAVN	Pickup	Book. Curve	EMSRb	Pickup	Book. Curve	232,231	224,225	*2.33%*	*-1.20%*
1.0	HBP	Pickup	Book. Curve	EMSRb	Pickup	Book. Curve	231,414	224,686	*1.97%*	*-1.00%*
1.0	EMSRb	Reg.	Book. Curve	EMSRb	Pickup	Book. Curve	227,859	227,191	*0.40%*	*0.11%*
1.0	VEMSRb	Reg.	Book. Curve	EMSRb	Pickup	Book. Curve	231,647	224,580	*2.07%*	*-1.05%*
1.0	Netbid	Reg.	Book. Curve	EMSRb	Pickup	Book. Curve	229,955	225,959	*1.32%*	*-0.44%*
1.0	DAVN	Reg.	Book. Curve	EMSRb	Pickup	Book. Curve	233,311	224,048	*2.80%*	*-1.28%*
1.0	HBP	Reg.	Book. Curve	EMSRb	Pickup	Book. Curve	232,400	224,561	*2.40%*	*-1.05%*
1.0	EMSRb	Reg.	Proj.	EMSRb	Pickup	Book. Curve	229,494	225,959	*1.12%*	*-0.44%*
1.0	VEMSRb	Reg.	Proj.	EMSRb	Pickup	Book. Curve	233,276	223,588	*2.79%*	*-1.48%*
1.0	Netbid	Reg.	Proj.	EMSRb	Pickup	Book. Curve	234,378	223,602	*3.27%*	*-1.48%*
1.0	DAVN	Reg.	Proj.	EMSRb	Pickup	Book. Curve	233,687	223,591	*2.97%*	*-1.48%*
1.0	HBP	Reg.	Proj.	EMSRb	Pickup	Book. Curve	234,287	223,251	*3.23%*	*-1.63%*
1.0	EMSRb	Pickup	Proj.	EMSRb	Pickup	Book. Curve	228,210	226,042	*0.55%*	*-0.40%*
1.0	VEMSRb	Pickup	Proj.	EMSRb	Pickup	Book. Curve	232,250	223,783	*2.33%*	*-1.40%*
1.0	Netbid	Pickup	Proj.	EMSRb	Pickup	Book. Curve	233,744	223,591	*2.99%*	*-1.48%*
1.0	DAVN	Pickup	Proj.	EMSRb	Pickup	Book. Curve	233,023	223,602	*2.67%*	*-1.48%*
1.0	HBP	Pickup	Proj.	EMSRb	Pickup	Book. Curve	233,279	223,358	*2.79%*	*-1.58%*
1.0	EMSRb	Pickup	Book. Curve	EMSRb	Reg.	Book. Curve	227,191	227,859	*0.10%*	*0.40%*
1.0	VEMSRb	Pickup	Book. Curve	EMSRb	Reg.	Book. Curve	230,425	225,800	*1.53%*	*-0.51%*
1.0	Netbid	Pickup	Book. Curve	EMSRb	Reg.	Book. Curve	228,913	227,191	*0.86%*	*0.11%*
1.0	DAVN	Pickup	Book Curve	EMSRb	Reg.	Book. Curve	232,230	225,128	*2.32%*	*-0.80%*
1.0	HBP	Pickup	Book. Curve	EMSRb	Reg.	Book. Curve	229,065	226,845	*0.93%*	*-0.05%*
1.0	EMSRb	Reg	Book. Curve	EMSRb	Reg.	Book. Curve	228,136	228,089	*0.52%*	*0.50%*
1.0	VEMSRb	Reg.	Book. Curve	EMSRb	Reg.	Book. Curve	231,669	225,501	*2.08%*	*-0.64%*
1.0	Netbid	Reg.	Book. Curve	EMSRb	Reg.	Book. Curve	230,009	226,969	*1.35%*	*0.01%*
1.0	DAVN	Reg.	Book. Curve	EMSRb	Reg.	Book. Curve	233,364	224,959	*2.82%*	*-0.88%*
1.0	HBP	Reg.	Book. Curve	EMSRb	Reg.	Book. Curve	232,470	225,486	*2.43%*	*-0.65%*
1.0	EMSRb	Reg	Proj	EMSRb	Reg.	Book. Curve	229,803	226,724	*1.26%*	*-0.10%*
1.0	VEMSRb	Reg.	Proj.	EMSRb	Reg.	Book. Curve	233,338	224,403	*2.81%*	*-1.12%*
1.0	Netbid	Reg.	Proj.	EMSRb	Reg.	Book. Curve	234,478	224,391	*3.32%*	*-1.13%*
1.0	DAVN	Reg.	Proj.	EMSRb	Reg.	Book. Curve	233,766	224,423	*3.00%*	*-1.11%*
1.0	HBP	Reg.	Proj.	EMSRb	Reg	Book. Curve	234,416	224,005	*3.29%*	*-1.30%*
1.0	EMSRb	Pickup	Proj.	EMSRb	Reg.	Book. Curve	228,467	226,834	*0.67%*	*-0.05%*
1.0	VEMSRb	Pickup	Proj.	EMSRb	Reg.	Book. Curve	232,271	224,623	*2.34%*	*-1.03%*
1.0	Netbid	Pickup	Proj.	EMSRb	Reg.	Book. Curve	233,816	224,393	*3.02%*	*-1.13%*
1.0	DAVN	Pickup	Proj.	EMSRb	Reg.	Book. Curve	233,112	224,423	*2.71%*	*-1.11%*
1.0	HBP	Pickup	Proj.	EMSRb	Reg.	Book. Curve	233,392	224,119	*2.84%*	*-1.25%*
1.0	EMSRb	Pickup	Book. Curve	EMSRb	Reg.	Proj.	225,959	229,494	*-0.44%*	*1.12%*
1.0	VEMSRb	Pickup	Book. Curve	EMSRb	Reg.	Proj.	229,691	227,163	*1.21%*	*0.09%*
1.0	Netbid	Pickup	Book. Curve	EMSRb	Reg.	Proj.	228,429	228,581	*0.65%*	*0.72%*
1.0	DAVN	Pickup	Book Curve	EMSRb	Reg.	Proj.	231,317	226,538	*1.92%*	*-0.18%*
1.0	HBP	Pickup	Book. Curve	EMSRb	Reg.	Proj.	228,336	228,119	*0.61%*	*0.51%*
1.0	EMSRb	Reg.	Book. Curve	EMSRb	Reg.	Proj.	226,724	229,803	*-0.10%*	*1.26%*
1.0	VEMSRb	Reg.	Book. Curve	EMSRb	Reg.	Proj.	230,696	226,897	*1.65%*	*-0.02%*
1.0	Netbid	Reg.	Book. Curve	EMSRb	Reg.	Proj.	229,436	228,423	*1.09%*	*0.65%*
1.0	DAVN	Reg.	Book. Curve	EMSRb	Reg.	Proj.	232,327	226,338	*2.37%*	*-0.27%*
1.0	HBP	Reg.	Book. Curve	EMSRb	Reg.	Proj.	228,096	228,489	*0.50%*	*0.68%*
1.0	EMSRb	Reg.	Proj.	EMSRb	Reg.	Proj.	227,977	227,956	*0.45%*	*0.44%*
1.0	VEMSRb	Reg.	Proj.	EMSRb	Reg.	Proj.	232,029	225,500	*2.24%*	*-0.64%*
1.0	Netbid	Reg.	Proj.	EMSRb	Reg.	Proj.	232,780	225,626	*2.57%*	*-0.58%*
1.0	DAVN	Reg.	Proj.	EMSRb	Reg.	Proj.	232,172	225,695	*2.30%*	*-0.55%*
1.0	HBP	Reg	Proj.	EMSRb	Reg.	Proj.	232,824	225,243	*2.59%*	*-0.75%*

1 0	EMSRb	Pickup	Proj.	EMSRb	Reg	Proj.	226,821	228,257	*-0.06%*	*0.58%*
1.0	VEMSRb	Pickup	Proj.	EMSRb	Reg.	Proj.	231,121	225,783	*1.84%*	*-0.52%*
1.0	Netbid	Pickup	Proj.	EMSRb	Reg.	Proj.	232,214	225,662	*2.32%*	*-0.57%*
1.0	DAVN	Pickup	Proj.	EMSRb	Reg.	Proj.	231,549	225,736	*2.02%*	*-0.54%*
1.0	HBP	Pickup	Proj.	EMSRb	Reg.	Proj.	232,169	225,631	*2.30%*	*-0.58%*
1.0	EMSRb	Pickup	Book. Curve	EMSRb	Pickup	Proj.	226,042	228,210	*-0.40%*	*0.55%*
1.0	VEMSRb	Pickup	Book. Curve	EMSRb	Pickup	Proj.	229,871	225,932	*1.29%*	*-0.45%*
1.0	Netbid	Pickup	Book. Curve	EMSRb	Pickup	Proj.	228,527	227,225	*0.69%*	*0.12%*
1.0	DAVN	Pickup	Book. Curve	EMSRb	Pickup	Proj.	231,526	225,289	*2.01%*	*-0.73%*
1.0	HBP	Pickup	Book. Curve	EMSRb	Pickup	Proj.	228,458	226,856	*0.66%*	*-0.04%*
1.0	EMSRb	Reg.	Book. Curve	EMSRb	Pickup	Proj.	226,834	228,467	*-0.05%*	*0.67%*
1.0	VEMSRb	Reg.	Book. Curve	EMSRb	Pickup	Proj.	230,904	225,688	*1.74%*	*-0.56%*
1.0	Netbid	Reg.	Book. Curve	EMSRb	Pickup	Proj.	229,524	227,090	*1.13%*	*0.06%*
1.0	DAVN	Reg.	Book. Curve	EMSRb	Pickup	Proj.	232,543	225,092	*2.46%*	*-0.82%*
1.0	HBP	Reg.	Book. Curve	EMSRb	Pickup	Proj.	228,178	227,202	*0.54%*	*0.11%*
1.0	EMSRb	Reg.	Proj.	EMSRb	Pickup	Proj.	228,257	226,821	*0.57%*	*-0.06%*
1.0	VEMSRb	Reg.	Proj.	EMSRb	Pickup	Proj.	232,320	224,386	*2.36%*	*-1.13%*
1.0	Netbid	Reg.	Proj.	EMSRb	Pickup	Proj.	233,086	224,545	*2.70%*	*-1.06%*
1.0	DAVN	Reg.	Proj.	EMSRb	Pickup	Proj.	232,507	224,568	*2.45%*	*-1.05%*
1.0	HBP	Reg.	Proj.	EMSRb	Pickup	Proj.	233,118	224,142	*2.72%*	*-1.24%*
1.0	EMSRb	Pickup	Proj.	EMSRb	Pickup	Proj.	227,061	226,975	*0.05%*	*0.01%*
1.0	VEMSRb	Pickup	Proj.	EMSRb	Pickup	Proj.	231,440	224,606	*1.98%*	*-1 03%*
1.0	Netbid	Pickup	Proj.	EMSRb	Pickup	Proj.	232,565	224,505	*2.47%*	*-1.08%*
1.0	DAVN	Pickup	Proj.	EMSRb	Pickup	Proj.	231,875	224,621	*2.17%*	*-1.03%*
1.0	HBP	Pickup	Proj.	EMSRb	Pickup	Proj.	232,249	224,306	*2.33%*	*-1.17%*

	-----Airline A-----			-----Airline B-----			*Revenues*		*% Increase compared to VEMSRb*	
DF	**Seat Optimizer**	**Forecast**	**Detruncation**	**Seat Optimizer**	**Forecast**	**Detruncation**	**Airline A**	**Airline B**	**Airline A**	**Airline B**
1.0	*EMSRb*	Pkup	Book Curve	VEMSRb	Pkup	Book Curve	224,877	230,425	*-1.40%*	*1.08%*
1.0	VEMSRb	*Pkup*	*Book Curve*	*VEMSRb*	*Pkup*	*Book Curve*	*228,079*	*227,964*	----	-----
1.0	Netbid	Pkup	Book Curve	VEMSRb	Pkup	Book Curve	225,898	229,238	*-0.96%*	*0.56%*
1.0	DAVN	Pkup	Book Curve	VEMSRb	Pkup	Book Curve	229,119	227,460	*0.46%*	*-0.22%*
1.0	HBP	Pkup	Book Curve	VEMSRb	Pkup	Book Curve	228,217	227,845	*0.06%*	*-0.05%*
1.0	EMSRb	Reg.	Book Curve	VEMSRb	Pkup	Book Curve	225,800	230,425	*-1.00%*	*1.08%*
1.0	VEMSRb	Reg.	Book Curve	VEMSRb	Pkup	Book Curve	229,069	227,792	*0 43%*	*-0.08%*
1.0	Netbid	Reg	Book Curve	VEMSRb	Pkup	Book Curve	227,087	229,031	*-0.43%*	*0.47%*
1.0	DAVN	Reg.	Book Curve	VEMSRb	Pkup	Book Curve	230,137	227,331	*0.90%*	*-0.28%*
1.0	HBP	Reg	Book Curve	VEMSRb	Pkup	Book Curve	229,067	227,781	*0.43%*	*-0.08%*
1.0	EMSRb	Reg.	Proj.	VEMSRb	Pkup	Book Curve	227,163	229,691	*-0.40%*	*0.76%*
1.0	VEMSRb	Reg.	Proj.	VEMSRb	Pkup	Book Curve	230,306	226,982	*0.98%*	*-0.43%*
1.0	Netbid	Reg.	Proj.	VEMSRb	Pkup	Book Curve	230,890	227,022	*1.23%*	*-0.41%*
1.0	DAVN	Reg.	Proj.	VEMSRb	Pkup	Book Curve	230,331	226,776	*0.99%*	*-0.52%*
1.0	HBP	Reg.	Proj.	VEMSRb	Pkup	Book Curve	230,856	226,660	*1.22%*	*-0.57%*
1.0	EMSRb	Pkup	Proj.	VEMSRb	Pkup	Book Curve	225,932	229,871	*-0.94%*	*0.84%*
1.0	VEMSRb	Pkup	Proj.	VEMSRb	Pkup	Book Curve	229,331	227,146	*0.55%*	*-0.36%*
1.0	Netbid	Pkup	Proj.	VEMSRb	Pkup	Book Curve	230,286	226,919	*0.97%*	*-0.46%*
1.0	DAVN	Pkup	Proj.	VEMSRb	Pkup	Book Curve	229,689	226,754	*0.71%*	*-0.53%*
1.0	HBP	Pkup	Proj.	VEMSRb	Pkup	Book Curve	229,902	226,749	*0.80%*	*-0.53%*
1.0	EMSRb	Pkup	Book Curve	VEMSRb	Reg.	Book Curve	224,580	231,647	*-1.53%*	*1.62%*
1.0	VEMSRb	Pkup	Book Curve	VEMSRb	Reg.	Book Curve	227,792	229,069	*-0.13%*	*0.48%*
1.0	Netbid	Pkup	Book Curve	VEMSRb	Reg.	Book Curve	225,762	230,478	*-1.02%*	*1.10%*
1.0	DAVN	Pkup	Book Curve	VEMSRb	Reg.	Book Curve	228,948	228,509	*0.38%*	*0.24%*
1.0	HBP	Pkup	Book Curve	VEMSRb	Reg.	Book Curve	227,996	229,008	*-0.04%*	*0.46%*
1.0	EMSRb	Reg.	Book Curve	VEMSRb	Reg.	Book Curve	225,501	231,669	*-1.13%*	*1.63%*
1.0	VEMSRb	Reg.	Book Curve	VEMSRb	Reg.	Book Curve	228,876	228,829	*0.35%*	*0.38%*
1.0	Netbid	Reg.	Book Curve	VEMSRb	Reg.	Book Curve	226,999	230,198	*-0.47%*	*0.98%*
1.0	DAVN	Reg.	Book Curve	VEMSRb	Reg.	Book Curve	230,012	228,355	*0.85%*	*0.17%*
1.0	HBP	Reg.	Book Curve	VEMSRb	Reg.	Book Curve	228,917	228,881	*0.37%*	*0.40%*

1.0	EMSRb	Reg.	Proj.	VEMSRb	Reg.	Book Curve	226,897	230,696	-0.52%	1.20%
1.0	VEMSRb	Reg.	Proj.	VEMSRb	Reg.	Book Curve	230,086	227,940	0.88%	-0.01%
1.0	Netbid	Reg.	Proj.	VEMSRb	Reg.	Book Curve	230,714	227,978	1.16%	0.01%
1.0	DAVN	Reg.	Proj.	VEMSRb	Reg.	Book Curve	230,156	227,791	0.91%	-0.08%
1.0	HBP	Reg.	Proj.	VEMSRb	Reg.	Book Curve	230,683	227,618	1.14%	-0.15%
1.0	EMSRb	Pkup	Proj.	VEMSRb	Reg.	Book Curve	225,688	230,904	-1.05%	1.29%
1.0	VEMSRb	Pkup	Proj.	VEMSRb	Reg.	Book Curve	229,115	228,109	0.45%	0.06%
1.0	Netbid	Pkup	Proj.	VEMSRb	Reg.	Book Curve	230,121	227,918	0.90%	-0.02%
1.0	DAVN	Pkup	Proj.	VEMSRb	Reg.	Book Curve	229,500	227,811	0.62%	-0.07%
1.0	HBP	Pkup	Proj.	VEMSRb	Reg.	Book Curve	229,733	227,721	0.73%	-0.11%
1.0	EMSRb	Pkup	Book Curve	VEMSRb	Reg.	Proj.	223,588	233,276	-1.97%	2.33%
1.0	VEMSRb	Pkup	Book Curve	VEMSRb	Reg.	Proj.	226,982	230,306	-0.48%	1.03%
1.0	Netbid	Pkup	Book Curve	VEMSRb	Reg.	Proj.	225,326	231,863	-1.21%	1.71%
1.0	DAVN	Pkup	Book Curve	VEMSRb	Reg.	Proj.	228,041	229,801	-0.02%	0.81%
1.0	HBP	Pkup	Book Curve	VEMSRb	Reg.	Proj.	227,304	230,219	-0.34%	0.99%
1.0	EMSRb	Reg	Book Curve	VEMSRb	Reg.	Proj.	224,403	233,338	-1.61%	2.36%
1.0	VEMSRb	Reg.	Book Curve	VEMSRb	Reg.	Proj.	227,940	230,086	-0.06%	0.93%
1.0	Netbid	Reg.	Book Curve	VEMSRb	Reg.	Proj.	226,427	231,645	-0.72%	1.61%
1.0	DAVN	Reg.	Book Curve	VEMSRb	Reg.	Proj.	229,033	229,624	0.42%	0.73%
1.0	HBP	Reg.	Book Curve	VEMSRb	Reg.	Proj.	228,069	230,122	0.00%	0.95%
1.0	EMSRb	Reg.	Proj.	VEMSRb	Reg.	Proj.	225,500	232,029	-1.13%	1.78%
1.0	VEMSRb	Reg	Proj	VEMSRb	Reg.	Proj.	228,953	228,920	0.38%	0.42%
1.0	Netbid	Reg.	Proj.	VEMSRb	Reg.	Proj.	229,431	229,172	0.59%	0.53%
1.0	DAVN	Reg.	Proj.	VEMSRb	Reg.	Proj.	228,792	229,102	0.31%	0.50%
1.0	HBP	Reg	Proj.	VEMSRb	Reg.	Proj.	229,425	228,751	0.59%	0.35%
1.0	EMSRb	Pkup	Proj.	VEMSRb	Reg.	Proj.	224,386	232,320	-1.62%	1.91%
1.0	VEMSRb	Pkup	Proj.	VEMSRb	Reg.	Proj.	228,019	229,173	-0.03%	0.53%
1.0	Netbid	Pkup	Proj.	VEMSRb	Reg	Proj.	228,815	229,164	0.32%	0.53%
1.0	DAVN	Pkup	Proj.	VEMSRb	Reg.	Proj.	228,110	229,154	0.01%	0.52%
1.0	HBP	Pkup	Proj	VEMSRb	Reg.	Proj.	228,537	228,935	0.20%	0.43%
1.0	EMSRb	Pkup	Book Curve	VEMSRb	Pkup	Proj.	223,783	232,250	-1.88%	1.88%
1.0	VEMSRb	Pkup	Book Curve	VEMSRb	Pkup	Proj.	227,146	229,331	-0.41%	0.60%
1.0	Netbid	Pkup	Book Curve	VEMSRb	Pkup	Proj.	225,389	230,768	-1.18%	1.23%
1.0	DAVN	Pkup	Book Curve	VEMSRb	Pkup	Proj.	228,203	228,824	0.05%	0.38%
1.0	HBP	Pkup	Book Curve	VEMSRb	Pkup	Proj	227,407	229,263	-0.29%	0.57%
1.0	EMSRb	Reg.	Book Curve	VEMSRb	Pkup	Proj.	224,623	232,271	-1.52%	1.89%
1.0	VEMSRb	Reg.	Book Curve	VEMSRb	Pkup	Proj.	228,109	229,115	0.01%	0.50%
1.0	Netbid	Reg.	Book Curve	VEMSRb	Pkup	Proj	226,479	230,569	-0.70%	1.14%
1.0	DAVN	Reg.	Book Curve	VEMSRb	Pkup	Proj.	229,150	228,675	0.47%	0.31%
1.0	HBP	Reg.	Book Curve	VEMSRb	Pkup	Proj.	228,152	229,177	0.03%	0.53%
1.0	EMSRb	Reg.	Proj.	VEMSRb	Pkup	Proj.	225,783	231,121	-1.01%	1.38%
1.0	VEMSRb	Reg	Proj.	VEMSRb	Pkup	Proj.	229,173	228,019	0.48%	0.02%
1.0	Netbid	Reg.	Proj.	VEMSRb	Pkup	Proj.	229,612	228,284	0.67%	0.14%
1.0	DAVN	Reg	Proj.	VEMSRb	Pkup	Proj.	228,972	228,185	0.39%	0.10%
1.0	HBP	Reg.	Proj.	VEMSRb	Pkup	Proj.	229,622	227,874	0.68%	-0.04%
1.0	EMSRb	Pkup	Proj.	VEMSRb	Pkup	Proj.	224,606	231,440	-1.52%	1.52%
1.0	VEMSRb	Pkup	Proj.	VEMSRb	Pkup	Proj.	228,310	228,175	0.10%	0.09%
1.0	Netbid	Pkup	Proj.	VEMSRb	Pkup	Proj.	229,075	228,258	0.44%	0.13%
1.0	DAVN	Pkup	Proj.	VEMSRb	Pkup	Proj.	228,388	228,217	0.14%	0.11%
1.0	HBP	Pkup	Proj	VEMSRb	Pkup	Proj.	228,802	227,988	0.32%	0.01%

	-----Airline A-----			------Airline B------			Revenues		% increase compared to Netbid	
DF	Seat Optimizer	Forecast	Detruncation	Seat Optimizer	Forecast	Detruncation	Airline A	Airline B	Airline A	Airline B
1.0	*EMSRb*	Pkup	Book Curve	Netbid	Pkup	Book Curve	226,151	228,849	-0.50%	0.73%
1.0	VEMSRb	Pkup	Book Curve	Netbid	Pkup	Book Curve	229,303	225,856	0.88%	-0.58%
1.0	*Netbid*	*Pkup*	*Book Curve*	*Netbid*	*Pkup*	*Book Curve*	*227,292*	*227,183*	-----	-----
1.0	DAVN	Pkup	Book Curve	Netbid	Pkup	Book Curve	230,771	225,585	1.53%	-0.70%
1.0	HBP	Pkup	Book Curve	Netbid	Pkup	Book Curve	229,795	225,803	1.10%	-0.61%

1.0	EMSRb	Reg.	Book Curve	Netbid	Pkup	Book Curve	227,191	228,913	-0.04%	0.76%
1.0	VEMSRb	Reg.	Book Curve	Netbid	Pkup	Book Curve	230,478	225,762	1.40%	-0.63%
1.0	Netbid	Reg.	Book Curve	Netbid	Pkup	Book Curve	228,422	227,161	0.50%	-0.01%
1.0	DAVN	Reg.	Book Curve	Netbid	Pkup	Book Curve	231,820	225,602	1.99%	-0.70%
1.0	HBP	Reg	Book Curve	Netbid	Pkup	Book Curve	230,786	225,883	1.54%	-0.57%
1.0	EMSRb	Reg.	Proj.	Netbid	Pkup	Book Curve	228,581	228,429	0.57%	0.55%
1.0	VEMSRb	Reg.	Proj.	Netbid	Pkup	Book Curve	231,863	225,326	2.01%	-0.82%
1.0	Netbid	Reg.	Proj.	Netbid	Pkup	Book Curve	232,562	225,735	2.32%	-0.64%
1.0	DAVN	Reg.	Proj.	Netbid	Pkup	Book Curve	232,015	225,668	2.08%	-0.67%
1.0	HBP	Reg.	Proj.	Netbid	Pkup	Book Curve	232,652	225,104	2.36%	-0.92%
1.0	EMSRb	Pkup	Proj.	Netbid	Pkup	Book Curve	227,225	228,527	-0.03%	0.59%
1.0	VEMSRb	Pkup	Proj.	Netbid	Pkup	Book Curve	230,768	225,389	1.53%	-0.79%
1.0	Netbid	Pkup	Proj.	Netbid	Pkup	Book Curve	231,894	225,662	2.02%	-0.67%
1.0	DAVN	Pkup	Proj.	Netbid	Pkup	Book Curve	231,339	225,577	1.78%	-0.71%
1.0	HBP	Pkup	Proj.	Netbid	Pkup	Book Curve	231,561	225,064	1.88%	-0.93%
1.0	EMSRb	Pkup	Book Curve	Netbid	Reg.	Book Curve	225,959	229,955	-0.59%	1.22%
1.0	VEMSRb	Pkup	Book Curve	Netbid	Reg.	Book Curve	229,031	227,087	0.77%	-0.04%
1.0	Netbid	Pkup	Book Curve	Netbid	Reg.	Book Curve	227,161	228,422	-0.06%	0.55%
1.0	DAVN	Pkup	Book Curve	Netbid	Reg.	Book Curve	230,602	226,645	1.46%	-0.24%
1.0	HBP	Pkup	Book Curve	Netbid	Reg	Book Curve	227,571	227,809	0.12%	0.28%
1.0	EMSRb	Reg.	Book Curve	Netbid	Reg.	Book Curve	226,969	230,009	-0.14%	1.24%
1.0	VEMSRb	Reg.	Book Curve	Netbid	Reg	Book Curve	230,198	226,999	1.28%	-0.08%
1.0	Netbid	Reg.	Book Curve	Netbid	Reg.	Book Curve	228,390	228,327	0.48%	0.50%
1.0	DAVN	Reg.	Book Curve	Netbid	Reg.	Book Curve	231,711	226,660	1.94%	-0.23%
1.0	HBP	Reg.	Book Curve	Netbid	Reg.	Book Curve	230,671	227,019	1.49%	-0.07%
1.0	EMSRb	Reg.	Proj.	Netbid	Reg.	Book Curve	228,423	229,436	0.50%	0.99%
1.0	VEMSRb	Reg.	Proj.	Netbid	Reg.	Book Curve	231,645	226,427	1.92%	-0.33%
1.0	Netbid	Reg.	Proj.	Netbid	Reg.	Book Curve	232,475	226,684	2.28%	-0.22%
1.0	DAVN	Reg	Proj.	Netbid	Reg	Book Curve	231,930	226,655	2.04%	-0.23%
1.0	HBP	Reg.	Proj.	Netbid	Reg.	Book Curve	232,523	226,140	2.30%	-0.46%
1.0	EMSRb	Pkup	Proj.	Netbid	Reg.	Book Curve	227,090	229,524	-0.09%	1.03%
1.0	VEMSRb	Pkup	Proj.	Netbid	Reg.	Book Curve	230,569	226,479	1.44%	-0.31%
1.0	Netbid	Pkup	Proj.	Netbid	Reg.	Book Curve	231,827	226,626	2.00%	-0.25%
1.0	DAVN	Pkup	Proj.	Netbid	Reg.	Book Curve	231,194	226,644	1.72%	-0.24%
1.0	HBP	Pkup	Proj	Netbid	Reg	Book Curve	231,517	226,012	1.86%	-0.52%
1.0	EMSRb	Pkup	Book Curve	Netbid	Reg.	Proj.	223,602	234,378	-1.62%	3.17%
1.0	VEMSRb	Pkup	Book Curve	Netbid	Reg	Proj.	227,022	230,890	-0.12%	1.63%
1.0	Netbid	Pkup	Book Curve	Netbid	Reg.	Proj.	225,735	232,562	-0.69%	2 37%
1.0	DAVN	Pkup	Book Curve	Netbid	Reg.	Proj.	228,525	230,368	0.54%	1.40%
1.0	HBP	Pkup	Book Curve	Netbid	Reg.	Proj.	225,774	231,913	-0.67%	2.08%
1.0	EMSRb	Reg.	Book Curve	Netbid	Reg.	Proj.	224,391	234,478	-1.28%	3.21%
1.0	VEMSRb	Reg.	Book Curve	Netbid	Reg	Proj.	227,978	230,714	0.30%	1.55%
1.0	Netbid	Reg	Book Curve	Netbid	Reg.	Proj.	226,684	232,475	-0.27%	2.33%
1.0	DAVN	Reg.	Book Curve	Netbid	Reg.	Proj.	229,465	230,245	0.96%	1.35%
1.0	HBP	Reg.	Book Curve	Netbid	Reg.	Proj.	225,632	232,371	-0.73%	2.28%
1.0	EMSRb	Reg.	Proj.	Netbid	Reg.	Proj.	225,626	232,780	-0.73%	2.46%
1.0	VEMSRb	Reg.	Proj.	Netbid	Reg.	Proj.	229,172	229,431	0.83%	0.99%
1.0	Netbid	Reg.	Proj.	Netbid	Reg.	Proj.	229,766	229,696	1.09%	1.11%
1.0	DAVN	Reg.	Proj.	Netbid	Reg.	Proj.	229,102	229,745	0.80%	1.13%
1.0	HBP	Reg.	Proj.	Netbid	Reg.	Proj.	229,688	229,121	1.05%	0.85%
1.0	EMSRb	Pkup	Proj.	Netbid	Reg.	Proj.	224,545	233,086	-1.21%	2.60%
1.0	VEMSRb	Pkup	Proj.	Netbid	Reg.	Proj.	228,284	229,612	0.44%	1.07%
1.0	Netbid	Pkup	Proj.	Netbid	Reg.	Proj.	229,193	229,710	0.84%	1.11%
1.0	DAVN	Pkup	Proj.	Netbid	Reg.	Proj.	228,413	229,849	0.49%	1.17%
1.0	HBP	Pkup	Proj.	Netbid	Reg.	Proj.	229,098	229,439	0.79%	0.99%

1.0	EMSRb	Pkup	Book Curve	Netbid	Pkup	Proj.	223,591	233,744	*-1.63%*	*2.89%*
1.0	VEMSRb	Pkup	Book Curve	Netbid	Pkup	Proj.	226,919	230,286	*-0.16%*	*1.37%*
1.0	Netbid	Pkup	Book Curve	Netbid	Pkup	Proj.	225,662	231,894	*-0.72%*	*2.07%*
1.0	DAVN	Pkup	Book Curve	Netbid	Pkup	Proj.	228,492	229,677	*0.53%*	*1.10%*
1.0	HBP	Pkup	Book Curve	Netbid	Pkup	Proj.	225,674	231,306	*-0.71%*	*1.81%*
1.0	EMSRb	Reg.	Book Curve	Netbid	Pkup	Proj.	224,393	233,816	*-1.28%*	*2.92%*
1.0	VEMSRb	Reg.	Book Curve	Netbid	Pkup	Proj.	227,918	230,121	*0.28%*	*1.29%*
1.0	Netbid	Reg.	Book Curve	Netbid	Pkup	Proj.	226,626	231,827	*-0.29%*	*2.04%*
1.0	DAVN	Reg.	Book Curve	Netbid	Pkup	Proj.	229,445	229,574	*0.95%*	*1.05%*
1.0	HBP	Reg.	Book Curve	Netbid	Pkup	Proj.	225,598	231,673	*-0.75%*	*1.98%*
1.0	EMSRb	Reg.	Proj.	Netbid	Pkup	Proj.	225,662	232,214	*-0.72%*	*2.21%*
1.0	VEMSRb	Reg.	Proj.	Netbid	Pkup	Proj.	229,164	228,815	*0.82%*	*0.72%*
1.0	Netbid	Reg.	Proj.	Netbid	Pkup	Proj.	229,710	229,193	*1.06%*	*0.88%*
1.0	DAVN	Reg.	Proj.	Netbid	Pkup	Proj.	229,115	229,124	*0.80%*	*0.85%*
1.0	HBP	Reg.	Proj.	Netbid	Pkup	Proj.	229,702	228,578	*1.06%*	*0.61%*
1.0	EMSRb	Pkup	Proj.	Netbid	Pkup	Proj.	224,505	232,565	*-1.23%*	*2.37%*
1.0	VEMSRb	Pkup	Proj.	Netbid	Pkup	Proj.	228,258	229,075	*0.43%*	*0.83%*
1.0	Netbid	Pkup	Proj.	Netbid	Pkup	Proj.	229,135	229,115	*0.81%*	*0.85%*
1.0	DAVN	Pkup	Proj.	Netbid	Pkup	Proj.	228,522	229,145	*0.54%*	*0.86%*
1.0	HBP	Pkup	Proj.	Netbid	Pkup	Proj.	228,927	228,583	*0.72%*	*0.62%*

	-----Airline A-----			-----Airline B-----			*Revenues*		*% Increase compared to DAVN*	
DF	**Seat Optimizer**	**Forecast**	**Detruncation**	**Seat Optimizer**	**Forecast**	**Detruncation**	**Airline A**	**Airline B**	**Airline A**	**Airline B**
1.0	*EMSRb*	Pkup	Book Curve	DAVN	Pkup	Book Curve	224,219	232,201	*-1.99%*	*1.52%*
1.0	VEMSRb	Pkup	Book Curve	DAVN	Pkup	Book Curve	227,514	229,080	*-0.55%*	*0.16%*
1.0	Netbid	Pkup	Book Curve	DAVN	Pkup	Book Curve	225,605	230,688	*-1.38%*	*0.86%*
1.0	*DAVN*	*Pkup*	*Book Curve*	*DAVN*	*Pkup*	*Book Curve*	*228,761*	*228,725*	----	----
1.0	HBP	Pkup	Book Curve	DAVN	Pkup	Book Curve	227,884	229,085	*-0.38%*	*0.16%*
1.0	EMSRb	Reg	Book Curve	DAVN	Pkup	Book Curve	225,128	232,230	*-1.59%*	*1.53%*
1.0	VEMSRb	Reg.	Book Curve	DAVN	Pkup	Book Curve	228,509	228,948	*-0.11%*	*0.10%*
1.0	Netbid	Reg.	Book Curve	DAVN	Pkup	Book Curve	226,645	230,602	*-0.92%*	*0.82%*
1.0	DAVN	Reg.	Book Curve	DAVN	Pkup	Book Curve	229,743	228,680	*0.43%*	*-0.02%*
1.0	HBP	Reg.	Book Curve	DAVN	Pkup	Book Curve	228,679	229,102	*-0.04%*	*0.16%*
1.0	EMSRb	Reg	Proj.	DAVN	Pkup	Book Curve	226,538	231,317	*-0.97%*	*1.13%*
1.0	VEMSRb	Reg.	Proj.	DAVN	Pkup	Book Curve	229,801	228,041	*0.45%*	*-0.30%*
1.0	Netbid	Reg	Proj.	DAVN	Pkup	Book Curve	230,368	228,525	*0.70%*	*-0.09%*
1.0	DAVN	Reg.	Proj.	DAVN	Pkup	Book Curve	229,556	228,464	*0.35%*	*-0.11%*
1 0	HBP	Reg.	Proj.	DAVN	Pkup	Book Curve	230,277	227,993	*0.66%*	*-0.32%*
1.0	EMSRb	Pkup	Proj.	DAVN	Pkup	Book Curve	225,289	231,526	*-1.52%*	*1.22%*
1.0	VEMSRb	Pkup	Proj.	DAVN	Pkup	Book Curve	228,824	228,203	*0.03%*	*-0.23%*
1.0	Netbid	Pkup	Proj.	DAVN	Pkup	Book Curve	229,677	228,492	*0.40%*	*-0.10%*
1.0	DAVN	Pkup	Proj.	DAVN	Pkup	Book Curve	228,958	228,454	*0.09%*	*-0.12%*
1.0	HBP	Pkup	Proj.	DAVN	Pkup	Book Curve	229,369	228,020	*0.27%*	*-0.31%*
1 0	EMSRb	Pkup	Book Curve	DAVN	Reg.	Book Curve	224,048	233,311	*-2.06%*	*2.01%*
1.0	VEMSRb	Pkup	Book Curve	DAVN	Reg.	Book Curve	227,331	230,137	*-0.63%*	*0.62%*
1.0	Netbid	Pkup	Book Curve	DAVN	Reg.	Book Curve	225,602	231,820	*-1.38%*	*1.35%*
1.0	DAVN	Pkup	Book Curve	DAVN	Reg.	Book Curve	228,680	229,743	*-0.04%*	*0.45%*
1.0	HBP	Pkup	Book Curve	DAVN	Reg.	Book Curve	227,672	230,202	*-0.48%*	*0.65%*
1.0	EMSRb	Reg.	Book Curve	DAVN	Reg.	Book Curve	224,959	233,364	*-1.66%*	*2.03%*
1.0	VEMSRb	Reg.	Book Curve	DAVN	Reg.	Book Curve	228,355	230,012	*-0.18%*	*0.56%*
1.0	Netbid	Reg.	Book Curve	DAVN	Reg.	Book Curve	226,660	231,711	*-0.92%*	*1.31%*
1.0	DAVN	Reg.	Book Curve	DAVN	Reg	Book Curve	229,674	229,671	*0.40%*	*0.41%*
1.0	HBP	Reg.	Book Curve	DAVN	Reg.	Book Curve	228,600	230,139	*-0.07%*	*0.62%*
1.0	EMSRb	Reg.	Proj.	DAVN	Reg.	Book Curve	226,338	232,327	*-1.06%*	*1.57%*
1.0	VEMSRb	Reg.	Proj.	DAVN	Reg.	Book Curve	229,624	229,033	*0.38%*	*0.13%*
1.0	Netbid	Reg.	Proj.	DAVN	Reg.	Book Curve	230,245	229,465	*0.65%*	*0.32%*
1.0	DAVN	Reg.	Proj.	DAVN	Reg.	Book Curve	229,489	229,433	*0.32%*	*0.31%*
1.0	HBP	Reg	Proj.	DAVN	Reg.	Book Curve	230,171	228,937	*0.62%*	*0.09%*

1.0	EMSRb	Pkup	Proj.	DAVN	Reg.	Book Curve	225,092	232,543	*-1.60%*	*1.67%*
1.0	VEMSRb	Pkup	Proj.	DAVN	Reg.	Book Curve	228,675	229,150	*-0.04%*	*0.19%*
1.0	Netbid	Pkup	Proj.	DAVN	Reg.	Book Curve	229,574	229,445	*0.36%*	*0.31%*
1.0	DAVN	Pkup	Proj.	DAVN	Reg.	Book Curve	228,835	229,427	*0.03%*	*0.31%*
1.0	HBP	Pkup	Proj.	DAVN	Reg.	Book Curve	229,250	228,997	*0.21%*	*0.12%*
1.0	EMSRb	Pkup	Book Curve	DAVN	Reg.	Proj.	223,591	233,687	*-2.26%*	*2.17%*
1.0	VEMSRb	Pkup	Book Curve	DAVN	Reg.	Proj.	226,776	230,331	*-0.87%*	*0.70%*
1.0	Netbid	Pkup	Book Curve	DAVN	Reg.	Proj.	225,668	232,015	*-1.35%*	*1.44%*
1.0	DAVN	Pkup	Book Curve	DAVN	Reg.	Proj.	228,464	229,556	*-0.13%*	*0.36%*
1.0	HBP	Pkup	Book Curve	DAVN	Reg.	Proj.	227,557	230,175	*-0.53%*	*0.63%*
1.0	EMSRb	Reg.	Book Curve	DAVN	Reg.	Proj.	224,423	233,766	*-1.90%*	*2.20%*
1.0	VEMSRb	Reg.	Book Curve	DAVN	Reg.	Proj.	227,791	230,156	*-0.42%*	*0.63%*
1.0	Netbid	Reg.	Book Curve	DAVN	Reg.	Proj.	226,655	231,930	*-0.92%*	*1.40%*
1.0	DAVN	Reg.	Book Curve	DAVN	Reg.	Proj.	229,433	229,489	*0.29%*	*0.33%*
1.0	HBP	Reg.	Book Curve	DAVN	Reg.	Proj.	228,288	230,146	*-0.21%*	*0.62%*
1.0	EMSRb	Reg.	Proj.	DAVN	Reg.	Proj.	225,695	232,172	*-1.34%*	*1.51%*
1.0	VEMSRb	Reg.	Proj.	DAVN	Reg.	Proj.	229,102	228,792	*0.15%*	*0.03%*
1.0	Netbid	Reg.	Proj.	DAVN	Reg	Proj.	229,745	229,102	*0.43%*	*0.16%*
1 0	DAVN	Reg	Proj.	DAVN	Reg.	Proj.	228,979	228,946	*0.10%*	*0 10%*
1.0	HBP	Reg.	Proj.	DAVN	Reg.	Proj.	229,666	228,311	*0.40%*	*-0.18%*
1.0	EMSRb	Pkup	Proj.	DAVN	Reg.	Proj.	224,568	232,507	*-1.83%*	*1.65%*
1.0	VEMSRb	Pkup	Proj.	DAVN	Reg.	Proj.	228,185	228,972	*-0.25%*	*0.11%*
1.0	Netbid	Pkup	Proj.	DAVN	Reg.	Proj.	229,124	229,115	*0.16%*	*0.17%*
1.0	DAVN	Pkup	Proj.	DAVN	Reg.	Proj.	228,301	228,984	*-0 20%*	*0.11%*
1.0	HBP	Pkup	Proj.	DAVN	Reg	Proj.	228,849	228,461	*0.04%*	*-0.12%*
1.0	EMSRb	Pkup	Book Curve	DAVN	Pkup	Proj	223,602	233,023	*-2 26%*	*1.88%*
1.0	VEMSRb	Pkup	Book Curve	DAVN	Pkup	Proj	226,754	229,689	*-0.88%*	*0.42%*
1.0	Netbid	Pkup	Book Curve	DAVN	Pkup	Proj.	225,577	231,339	*-1.39%*	*1.14%*
1.0	DAVN	Pkup	Book Curve	DAVN	Pkup	Proj.	228,454	228,958	*-0.13%*	*0.10%*
1.0	HBP	Pkup	Book Curve	DAVN	Pkup	Proj.	227,523	229,502	*-0.54%*	*0.34%*
1.0	EMSRb	Reg.	Book Curve	DAVN	Pkup	Proj.	224,423	233,112	*-1.90%*	*1.92%*
1.0	VEMSRb	Reg.	Book Curve	DAVN	Pkup	Proj.	227,811	229,500	*-0.42%*	*0.34%*
1.0	Netbid	Reg.	Book Curve	DAVN	Pkup	Proj.	226,644	231,194	*-0.93%*	*1.08%*
1 0	DAVN	Reg.	Book Curve	DAVN	Pkup	Proj.	229,427	228,835	*0 29%*	*0 05%*
1.0	HBP	Reg	Book Curve	DAVN	Pkup	Proj.	228,233	229,526	*-0.23%*	*0.35%*
1.0	EMSRb	Reg.	Proj.	DAVN	Pkup	Proj.	225,736	231,549	*-1.32%*	*1.23%*
1.0	VEMSRb	Reg.	Proj.	DAVN	Pkup	Proj.	229,154	228,110	*0.17%*	*-0.27%*
1 0	Netbid	Reg.	Proj.	DAVN	Pkup	Proj.	229,849	228,413	*0 48%*	*-0.14%*
1 0	DAVN	Reg.	Proj.	DAVN	Pkup	Proj.	228,984	228,301	*0.10%*	*-0.19%*
1.0	HBP	Reg.	Proj.	DAVN	Pkup	Proj.	229,732	227,674	*0.42%*	*-0 46%*
1.0	EMSRb	Pkup	Proj.	DAVN	Pkup	Proj.	224,621	231,875	*-1.81%*	*1.38%*
1.0	VEMSRb	Pkup	Proj.	DAVN	Pkup	Proj.	228,217	228,388	*-0.24%*	*-0.15%*
1.0	Netbid	Pkup	Proj.	DAVN	Pkup	Proj.	229,145	228,522	*0.17%*	*-0.09%*
1.0	DAVN	Pkup	Proj.	DAVN	Pkup	Proj.	228,381	228,272	*-0.17%*	*-0 20%*
1.0	HBP	Pkup	Proj.	DAVN	Pkup	Proj.	228,922	227,718	*0.07%*	*-0.44%*

	-----Airline A-----			-----Airline B-----			*Revenues*		***% increase compared to HBP***	
DF	**Seat Optimizer**	**Forecast**	**Detruncation**	**Seat Optimizer**	**Forecast**	**Detruncation**	**Airline A**	**Airline B**	**Airline A**	**Airline B**
1.0	*EMSRb*	Pkup	Book Curve	HBP	Pkup	Book Curve	225,810	229,101	*-1.09%*	*0.38%*
1.0	VEMSRb	Pkup	Book Curve	HBP	Pkup	Book Curve	227,919	228,157	*-0.17%*	*-0.03%*
1.0	Netbid	Pkup	Book Curve	HBP	Pkup	Book Curve	226,521	227,787	*-0.78%*	*-0.19%*
1.0	DAVN	Pkup	Book Curve	HBP	Pkup	Book Curve	229,195	227,750	*0.39%*	*-0.21%*
1.0	*HBP*	*Pkup*	*Book Curve*	*HBP*	*Pkup*	*Book Curve*	*228,305*	*228,226*	-----	-----

1.0	EMSRb	Reg.	Book Curve	HBP	Pkup	Book Curve	226,845	229,065	-0.64%	0.37%
1.0	VEMSRb	Reg.	Book Curve	HBP	Pkup	Book Curve	229,008	227,996	0.31%	-0.10%
1.0	Netbid	Reg.	Book Curve	HBP	Pkup	Book Curve	227,809	227,571	-0.22%	-0.29%
1.0	DAVN	Reg.	Book Curve	HBP	Pkup	Book Curve	230,202	227,672	0.83%	-0.24%
1.0	HBP	Reg.	Book Curve	HBP	Pkup	Book Curve	229,203	228,225	0.39%	0.00%
1.0	EMSRb	Reg.	Proj.	HBP	Pkup	Book Curve	228,119	228,336	-0.08%	0.05%
1.0	VEMSRb	Reg.	Proj.	HBP	Pkup	Book Curve	230,219	227,304	0.84%	-0.40%
1.0	Netbid	Reg.	Proj.	HBP	Pkup	Book Curve	231,913	225,774	1.58%	-1.07%
1.0	DAVN	Reg.	Proj.	HBP	Pkup	Book Curve	230,175	227,557	0.82%	-0.29%
1.0	HBP	Reg.	Proj.	HBP	Pkup	Book Curve	230,765	227,244	1.08%	-0.43%
1.0	EMSRb	Pkup	Proj.	HBP	Pkup	Book Curve	226,856	228,458	-0.63%	0.10%
1.0	VEMSRb	Pkup	Proj.	HBP	Pkup	Book Curve	229,263	227,407	0.42%	-0.36%
1.0	Netbid	Pkup	Proj.	HBP	Pkup	Book Curve	231,306	225,674	1.31%	-1.12%
1.0	DAVN	Pkup	Proj.	HBP	Pkup	Book Curve	229,502	227,523	0.52%	-0.31%
1.0	HBP	Pkup	Proj.	HBP	Pkup	Book Curve	229,848	227,221	0.68%	-0.44%
1.0	EMSRb	Pkup	Book Curve	HBP	Reg.	Book Curve	224,561	232,400	-1.64%	1.83%
1.0	VEMSRb	Pkup	Book Curve	HBP	Reg.	Book Curve	227,781	229,067	-0.23%	0.37%
1.0	Netbid	Pkup	Book Curve	HBP	Reg	Book Curve	225,883	230,786	-1.06%	1.12%
1.0	DAVN	Pkup	Book Curve	HBP	Reg.	Book Curve	229,102	228,679	0.35%	0.20%
1.0	HBP	Pkup	Book Curve	HBP	Reg.	Book Curve	228,225	229,203	-0.04%	0.43%
1.0	EMSRb	Reg.	Book Curve	HBP	Reg.	Book Curve	225,486	232,470	-1.23%	1.86%
1.0	VEMSRb	Reg	Book Curve	HBP	Reg.	Book Curve	228,881	228,917	0.25%	0.30%
1.0	Netbid	Reg	Book Curve	HBP	Reg.	Book Curve	227,019	230,671	-0.56%	1.07%
1.0	DAVN	Reg.	Book Curve	HBP	Reg.	Book Curve	230,139	228,600	0.80%	0.16%
1.0	HBP	Reg.	Book Curve	HBP	Reg.	Book Curve	229,178	229,177	0.38%	0.42%
1.0	EMSRb	Reg.	Proj.	HBP	Reg.	Book Curve	228,489	228,096	0.08%	-0.06%
1.0	VEMSRb	Reg	Proj.	HBP	Reg.	Book Curve	230,122	228,069	0.80%	-0.07%
1.0	Netbid	Reg.	Proj.	HBP	Reg.	Book Curve	232,371	225,632	1.78%	-1.14%
1.0	DAVN	Reg.	Proj.	HBP	Reg.	Book Curve	230,146	228,288	0.81%	0.03%
1.0	HBP	Reg	Proj.	HBP	Reg	Book Curve	230,728	227,993	1.06%	-0.10%
1.0	EMSRb	Pkup	Proj	HBP	Reg.	Book Curve	227,202	228,178	-0.48%	-0.02%
1.0	VEMSRb	Pkup	Proj.	HBP	Reg.	Book Curve	229,177	228,152	0.38%	-0.03%
1.0	Netbid	Pkup	Proj.	HBP	Reg.	Book Curve	231,673	225,598	1.48%	-1.15%
1.0	DAVN	Pkup	Proj.	HBP	Reg.	Book Curve	229,526	228,233	0.53%	0.00%
1.0	HBP	Pkup	Proj.	HBP	Reg.	Book Curve	229,815	227,985	0.66%	-0.11%
1.0	EMSRb	Pkup	Book Curve	HBP	Reg.	Proj.	223,251	234,287	-2.21%	2.66%
1.0	VEMSRb	Pkup	Book Curve	HBP	Reg.	Proj.	226,660	230,856	-0.72%	1.15%
1.0	Netbid	Pkup	Book Curve	HBP	Reg.	Proj.	225,104	232,652	-1.40%	1.94%
1.0	DAVN	Pkup	Book Curve	HBP	Reg.	Proj.	227,993	230,277	-0.14%	0.90%
1.0	HBP	Pkup	Book Curve	HBP	Reg.	Proj.	227,244	230,765	-0.46%	1.11%
1.0	EMSRb	Reg.	Book Curve	HBP	Reg.	Proj.	224,005	234,416	-1.88%	2.71%
1.0	VEMSRb	Reg.	Book Curve	HBP	Reg.	Proj.	227,618	230,683	-0.30%	1.08%
1.0	Netbid	Reg.	Book Curve	HBP	Reg.	Proj.	226,140	232,523	-0.95%	1.88%
1.0	DAVN	Reg.	Book Curve	HBP	Reg.	Proj.	228,937	230,171	0.28%	0.85%
1.0	HBP	Reg.	Book Curve	HBP	Reg.	Proj.	227,993	230,728	-0.14%	1.10%
1.0	EMSRb	Reg.	Proj.	HBP	Reg.	Proj.	225,243	232,824	-1.34%	2.01%
1.0	VEMSRb	Reg.	Proj.	HBP	Reg.	Proj.	228,751	229,425	0.20%	0.53%
1.0	Netbid	Reg.	Proj.	HBP	Reg.	Proj.	229,121	229,688	0.36%	0.64%
1.0	DAVN	Reg.	Proj.	HBP	Reg.	Proj.	228,311	229,666	0.00%	0.63%
1.0	HBP	Reg.	Proj.	HBP	Reg.	Proj.	229,149	229,129	0.37%	0.40%
1.0	EMSRb	Pkup	Proj.	HBP	Reg.	Proj.	224,142	233,118	-1.82%	2.14%
1.0	VEMSRb	Pkup	Proj.	HBP	Reg.	Proj.	227,874	229,622	-0.19%	0.61%
1.0	Netbid	Pkup	Proj.	HBP	Reg.	Proj.	228,578	229,702	0.12%	0.65%
1.0	DAVN	Pkup	Proj.	HBP	Reg.	Proj.	227,674	229,732	-0.28%	0.66%
1.0	HBP	Pkup	Proj.	HBP	Reg.	Proj.	228,305	229,269	0.00%	0.46%

1.0	EMSRb	Pkup	Book Curve	HBP	Pkup	Proj.	223,358	233,279	-2.17%	2.21%
1.0	VEMSRb	Pkup	Book Curve	HBP	Pkup	Proj.	226,749	229,902	-0.68%	0.73%
1.0	Netbid	Pkup	Book Curve	HBP	Pkup	Proj.	225,064	231,561	-1.42%	1.46%
1.0	DAVN	Pkup	Book Curve	HBP	Pkup	Proj.	228,020	229,369	-0.12%	0.50%
1.0	HBP	Pkup	Book Curve	HBP	Pkup	Proj.	227,221	229,848	-0.47%	0.71%
1.0	EMSRb	Reg.	Book Curve	HBP	Pkup	Proj.	224,119	233,392	-1.83%	2.26%
1.0	VEMSRb	Reg.	Book Curve	HBP	Pkup	Proj.	227,721	229,733	-0.26%	0.66%
1.0	Netbid	Reg.	Book Curve	HBP	Pkup	Proj.	226,012	231,517	-1.00%	1.44%
1.0	DAVN	Reg.	Book Curve	HBP	Pkup	Proj.	228,997	229,250	0.30%	0.45%
1.0	HBP	Reg.	Book Curve	HBP	Pkup	Proj.	227,985	229,815	-0.14%	0.70%
1.0	EMSRb	Reg.	Proj.	HBP	Pkup	Proj.	225,631	232,169	-1.17%	1.73%
1.0	VEMSRb	Reg.	Proj.	HBP	Pkup	Proj.	228,935	228,537	0.28%	0.14%
1.0	Netbid	Reg.	Proj.	HBP	Pkup	Proj.	229,439	229,098	0.50%	0.38%
1.0	DAVN	Reg.	Proj.	HBP	Pkup	Proj.	228,461	228,849	0.07%	0.27%
1.0	HBP	Reg.	Proj.	HBP	Pkup	Proj.	229,269	228,305	0.42%	0.03%
1.0	EMSRb	Pkup	Proj.	HBP	Pkup	Proj.	224,306	232,249	-1.75%	1.76%
1.0	VEMSRb	Pkup	Proj.	HBP	Pkup	Proj.	227,988	228,802	-0.14%	0.25%
1.0	Netbid	Pkup	Proj.	HBP	Pkup	Proj.	228,583	228,927	0.12%	0.31%
1.0	DAVN	Pkup	Proj.	HBP	Pkup	Proj.	227,718	228,922	-0.26%	0.30%
1.0	HBP	Pkup	Proj.	HBP	Pkup	Proj.	228,487	228,356	0.08%	0 06%

Bibliography

Belobaba, Peter P. Network Seat Inventory Control Without Network Optimization. Presented at the AGIFORS Yield Management Study Group, Hong Kong, January 17-19, 1994.

Belobaba, Peter P. "The Evolution of Airline Yield Management: Fare Class to Origin-Destination Seat Inventory Control." Handbook of Airline Marketing, 1st Ed. Aviation Week Group, 1998, pp. 285-302.

Belobaba, Peter P. Optimal vs. Heuristic Methods for Nested Seat Allocation. Presented at the AGIFORS Yield Management Study Group, Brussels, Belgium: May 4, 1992.

Belobaba, Peter P. "The Revenue Enhancement Potential of Airline Yield Management Systems." ASTAIR 1992 International Conference and Exhibition Proceedings. London, England: November 9-11, 1992.

Belobaba, Peter P and Craig A. Hopperstad. Competitive Impacts of Origin-Destination Control. Presented at INFORMS, San Diego, CA: May 4, 1997.

Belobaba, Peter P. Air Travel Demand and Airline Seat Inventory Management. Cambridge, MA: Massachusetts Institute of Technology, July, 1987.

Belobaba, Peter P. Airline O-D Seat Inventory Control Without Network Optimization. Cambridge, MA: Massachusetts Institute of Technology, June, 1995.

Bodily, Samuel E. and Lawrence R. Weatherford. "A Taxonomy and Research Overview of Perishable-Asset Revenue Management: Yield Management, Overbooking, and Pricing." OR Forum. Vol. 40, September-October, 1992, pp. 831-844.

Bratu, Stephane. Network Value Concept in Airline Revenue Management. Cambridge, MA: Massachusetts Institute of Technology, June, 1998.

Decision Window Path Preference Methodology Time Mode. Seattle, WA: Boeing Commercial Airplane Group, November, 1993.

Hopperstad, Craig A. "Passenger O/D Simulator--PODS v.6." 36th AGIFORS Annual Symposium Proceedings, Atlanta, GA: November 6-10, 1996.

Hopperstad, Craig A. "Passenger O/D Simulator--PODS v.6." Boeing Company Memo. Renton, WA: October, 1996.

Hopperstad, Craig A. Projection Detruncation. Boeing Company Memo. Renton, WA: October 30, 1997.

Hopperstad, Craig A. PODS: Modeling Update. Presented at the AGIFORS Yield Management Study Group, Montreal, Canada: May 14-16, 1997.

Lee, Alex Y. Investigation of Competitive Impacts of Origin-Destination Control Using PODS. Cambridge, MA: Massachusetts Institute of Technology, June, 1998.

Lee, Alex Y. PODS Simulation Results: Scenario X--Ten City Network. Presented at the 6th PODS Summit, Seattle, WA, April, 1998.

L'Heureux, Ed. "A New Twist in Forecasting Short-Term Passenger Pickup." 26th AGIFORS Annual Symposium Proceedings, Bowness-on-Windermere, England: October 19-24, 1986.

Pindyck, Robert S. & Rubenfeld, D. Econometric Models and Econometric Forecasts, 4th Ed. Boston, MA: McGraw-Hill, 1998.

Simpson, Robert W. Using Network Flow Techniques to Find Shadow Prices for Market Demands and Seat Inventory Control. Cambridge, MA: Massachusetts Institute of Technology, Flight Transportation Laboratory Memorandum M89-1, January, 1989.

Skwarek, Daniel K. Competitive Impacts of Yield Management System Components: Forecasting and Sell-Up Models. Cambridge, MA: Massachusetts Institute of Technology, August, 1996.

Smith, Barry C., et al. "Yield Management at American Airlines." Interfaces. January, 1992.

Swan, Bill. "Bid Pricing: Where it Came From, Where it May Go." 34th AGIFORS Annual Symposium Proceedings, Zandvoort, Netherlands: October 2-6, 1994.

Wei, Julie. Airline O-D Control Using Network Displacement Costs. Cambridge, MA: Massachusetts Institute of Technology, June, 1997.

Wickham, Richard R. Evaluation of Forecasting Techniques for Short-Term Demand of Air Transportation. Cambridge, MA: Massachusetts Institute of Technology, June, 1995.

Williamson, Elizabeth L. Airline Network Seat Inventory Control: Methodologies and Revenue Impacts. Cambridge, MA: Massachusetts Institute of Technology, June, 1992.

Wilson, John L. The Value of Revenue Management Innovation in a Competitive Airline Industry. Cambridge, MA: Massachusetts Institute of Technology, May, 1995.