Improved Forecast Accuracy in Airline Revenue Management by Unconstraining Demand Estimates from Censored Data

by

Richard H. Zeni

ISBN: 1-58112-141-5

DISSERTATION.COM

USA • 2001

Improved Forecast Accuracy in Airline Revenue Management by Unconstraining Demand Estimates from Censored Data

Dissertation.com
USA • 2001

ISBN: 1-58112-141-5

www.Dissertation.com/library/1121415a.htm

IMPROVED FORECAST ACCURACY IN REVENUE MANAGEMENT BY UNCONSTRAINING DEMAND ESTIMATES FROM CENSORED DATA

By Richard H. Zeni

A dissertation submitted to the

Graduate School-Newark

Rutgers, The State University of New Jersey

in partial fulfillment of requirements

for the degree of

Doctor of Philosophy

Ph.D. in Management Program

Written under the direction of

Professor Kenneth D. Lawrence

Newark, New Jersey

October, 2001

ABSTRACT OF THE DISSERTATION

Improved Forecast Accuracy in Revenue Management by

Unconstraining Demand Estimates from Censored Data

By Richard H. Zeni

Dissertation director: Professor Kenneth D. Lawrence

Accurate forecasts are crucial to a revenue management system. Poor estimates of demand lead to inadequate inventory controls and sub-optimal revenue performance. Forecasting for airline revenue management systems is inherently difficult. Competitive actions, seasonal factors, the economic environment, and constant fare changes are a few of the hurdles that must be overcome. In addition, the fact that most of the historical demand data is censored further complicates the problem. This dissertation examines the challenge of forecasting for an airline revenue management system in the presence of censored demand data.

The number of seats an airline can sell on a flight is determined by the booking limits set by the revenue management system. An airline continues to accept reservations in a fare class until the booking limit is reached. At that point, the airline stops selling seats in that fare class-It also stops collecting valuable data. Demand for travel in that fare class may exceed the booking limit, but the data does not reflect this. So the data is censored or “constrained” at the booking limit.

While some models exist that produce unbiased forecasts from censored data, it is preferable to “unconstrain” the censored observations so that they represent true demand.

Then, the forecasting model may be chosen based on the structure of the problem rather than the nature of the data. This dissertation analyzed the improvement in forecast accuracy that results from estimating demand by unconstraining the censored data.

Little research has been done on unconstraining censored data for revenue management systems. Airlines tend to either ignore the problem or use very simple ad hoc methods to deal with it. A literature review explores the current methods for unconstraining censored data. Also, practices borrowed from areas outside of revenue management are adapted to this application. For example, the Expectation-Maximization (EM) and other imputation methods were investigated. These methods are evaluated and tested using simulation and actual airline data. An extension to the EM algorithm that results in a 41% improvement in forecast accuracy is presented.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

Accurate forecasts are crucial to a revenue management system. Poor estimates of demand lead to incorrect decision making and less-than-optimal revenue performance. Revenue management was invented by the U.S. airlines in the 1980's in response to a newly deregulated industry and to the increased competition that was created. Since then, it has been adopted by a variety of industries, and the list is constantly growing. But the basic concepts have been around for quite a long time. Consider the grocer who has a supply of perishable fruit. He must price the fruit so that he maximizes his revenues and avoids having the fruit spoil. If he prices it too low, he will sell out too soon and will need to turn away customers when his stock runs out. He would have been better off setting a higher price. If he prices the fruit too high, he will not sell his entire stock and will be forced to dispose of some fruit and take a loss on that inventory.

The grocer makes his pricing decision based on his *forecast* of demand for the fruit. He must predict the quantity of fruit demanded at different price levels. As time goes by and the fruit approaches its final sale date, the grocer will reforecast demand and adjust his price accordingly. For example, he might lower his price if he sees that the fruit is not selling as fast as he thought and is about to spoil. Forecasting the demand for his product is more difficult when he runs out of fruit. It is difficult for him to estimate how much fruit he would have sold if he had more in stock. In this sense, his data is *censored.*

The example of the grocer illustrates a low-tech version of revenue management and of the challenge of forecasting demand with censored data. A more sophisticated application of revenue management is practiced by the airlines. They amass large numbers of reservations booking data, and they use it to produce forecasts of demand for airline travel at different price levels. Then they use complicated optimization models to set seat inventory levels for different products. They repeat the forecasting and optimization until the flight departs and the leftover seats spoil, just as the perishable fruit spoils. Hotels use revenue management to overbook rooms and discount rates for non-peak periods.[1] Rental car companies, theaters, and radio stations all use revenue management to manage their perishable inventory.[2]

Revenue management was developed by the airlines to improve revenue performance in the face of increasing competition. It was obvious to the airlines that passengers could be divided into two broad categories, based on their travel behavior and their sensitivity to prices. There were business travelers and leisure travelers. Business passengers tended to make their travel arrangements close to their departure date and stay at their destination for only a short time. There was usually little flexibility in their plans, and they were willing to pay higher prices for tickets. Leisure travelers, on the other hand, booked their flights well in advance of their travel date. They stayed longer at their destinations and had much more flexibility in their plans. They would often decide not to travel rather than pay high fares, and flights often departed with empty seats.

The challenge to the innovators of revenue management was to devise a plan that would make the empty seats available at the lower fair while preventing passengers who were willing to pay the higher fare from buying low-fare seats. Since low-fare passengers typically book in before higher fare customers, the revenue management system must forecast how many business passengers will want to book on a flight. Then it must set aside or "protect" these seats so that they will be available when the business passengers request them.

Accurate forecasts of passenger demand are crucial. If the forecast for business passengers is too high, then too many seats will be protected for these passengers. The flight will depart with empty seats that could have been sold to leisure passengers. If the forecast for business passengers is too low, then too few seats will be protected. Seats will be sold to leisure passengers that could have been sold to higher-fare business passengers.

Forecasting demand accurately is inherently difficult since the historical data upon which forecasts are based often do not reflect the true demand. Once an airline stops selling tickets at a particular fare, due to the limits set by the revenue management system, it also stops collecting data. The airline may receive many more requests for a particular fare, but these requests are not recorded – the data is censored and does not represent true demand.

When censored data is used to represent historical demand, it often results in forecasts with a negative bias. If a revenue management system uses these biased forecasts, then the resulting inventory controls will tend to save too few seats for high-fare passengers. Seats that could be sold to high-fare passengers may be sold to low-fare passengers, and revenue will be lost. It has been estimated that up to 3.0% of potential revenue may be lost if the forecasts used by a revenue management system have a negative bias.[3]

Therefore, some attempt should be made to transform the censored data into more accurate estimates of actual historical demand. Various methods exist that take the observations that have been "constrained" and "unconstrain" them so they represent the actual demand. These methods range from the simplistic, such as discarding all censored observations, to the complex, such as calculating the expected value of the true demand via the Expectation-Maximization algorithm. However, little research has been done to determine which methods work best. Airlines tend to use the simple heuristic methods. While they recognize that these heuristics are not adequate, they are hesitant to invest in other techniques due to the lack of evidence that alternative "unconstraining" will produce more accurate forecasts.

Airlines use complex revenue management systems to determine the number of seats to make available at different fares. First, a forecast of demand is produced from censored data. Based on that forecast, booking limits at the various fare levels are set. Reservation requests are then either accepted or denied based on the booking limit. The bookings that are accepted become the historical data for the next forecast. The bookings that are denied are not recorded, hence the censored data. The censored data is then unconstrained so that it represents the true demand, and the process begins again. This feedback loop is illustrated in Figure 1.1.

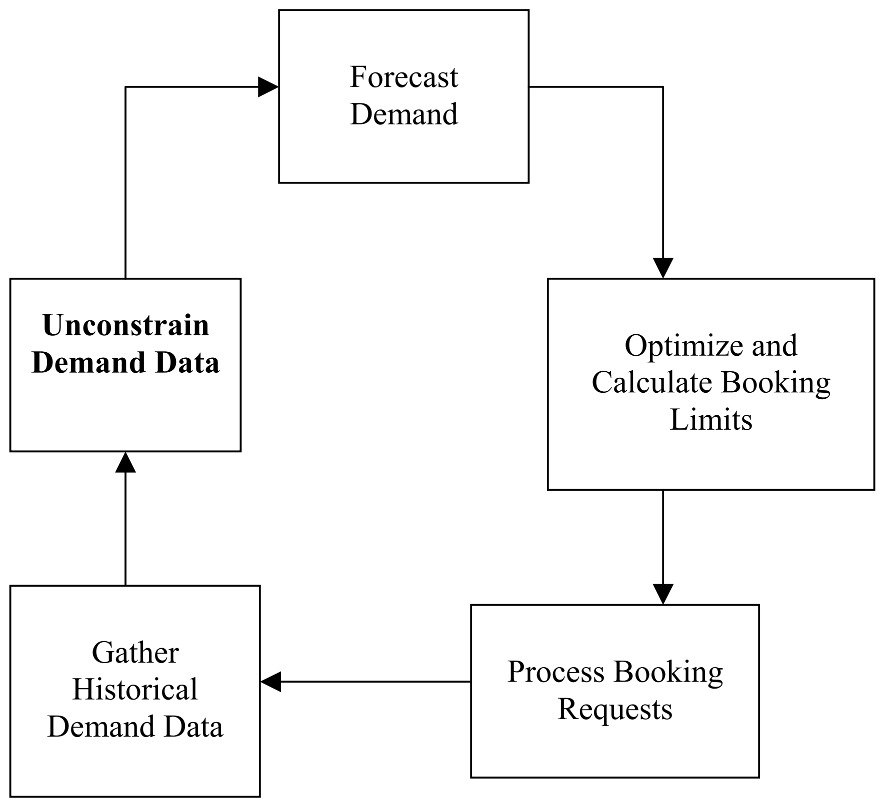

Figure 1.1: Revenue Management Process

1.1 Goal of Dissertation

This research project examines the challenge of forecasting in the presence of censored data. The goal of this dissertation is to improve forecast accuracy by unconstraining the constrained observations. Each of the unconstraining methods are examined and compared to determine which methods produce the best estimate of true demand. The degree of improvement in forecast accuracy that can be gained by unconstraining data were investigated. The costs of using this data in a revenue management system are investigated. While this dissertation will focus on forecasting for airline revenue management, the methods discussed and developed here easily generalize to other industries.

1.2 Structure of Dissertation

The remainder of this paper is divided into five chapters. The following sections in this chapter are a review of what revenue management is and how airlines and other transportation companies use it in an attempt to maximize revenue. The different levels at which revenue management can be applied are introduced and illustrated. The various optimization methods are surveyed.

Chapter 2 contains a literature review of forecasting for revenue management, and it surveys the current practices. Also, the discussion of censored data begins here. The challenges that this data present to forecasting are examined, and the cost of ignoring the problem is investigated. The literature is reviewed to present an overview of the available methods for dealing with the censored data problem. Simple ad-hoc and more complex statistical methods are discussed. The strengths and weaknesses of each are examined, and applications outside of the transportation industry are discussed.

Prior to analyzing methods for unconstraining censored data, a reliable data set needed be built to simulate censored and uncensored data. In Chapter 3, a new technique developed as part of this research project is presented. This method transforms uncensored data into data with censored observations.

Chapter 4 contains a comprehensive analysis and comparison of the available unconstraining methods. Each method is applied to a sample data set. The improvement in forecast accuracy from each of these techniques is evaluated. An extension to the EM algorithm developed as part of this research project is presented.

Chapter 5 concludes the dissertation. The research findings are summarized, and future research directions are outlined. Application of the analyzed methods is discussed with respect to practical considerations.

1.3 Revenue Management Defined

There are several definitions of revenue management (also referred to as yield management) in the literature. American Airlines (1987) defined the goal of yield management as "to maximize passenger revenue by selling the right seats to the right customers at the right time."[4] Pfiefer (1989) described airline yield management as the "process by which discount fares are allocated to scheduled flights for the purposes of balancing demand and increasing revenues."[5] From the hotel industry's perspective it has been defined as "charging a different rate for the same service to a different individual" [6] and "controlling the tradeoff between average rate and occupancy."[7]

Weatherford and Bodily (1992) have concluded from the above definitions that the term yield management is too limited in describing the broad class of revenue management approaches.[8] After analyzing situations in which yield management was used, they concluded that these situations had the following characteristics in common:

1. **There is one date on which the product or service becomes available and another after which it is either not available or it spoils.** The product cannot be stored for significant periods of time-It eventually perishes. In the grocery store example, the fruit would spoil.

2. **There is a fixed number of units.** Capacity cannot be changed in the short term. In the hotel example, there are only so many rooms that may be sold at a given property location.
3. **There is the possibility of segmenting price-sensitive customers.** In the airline example, vacation travelers are much more sensitive to price than business travelers.

Weatherford and Bodily proposed the term perishable-asset revenue management (PARM) to define this class of problems and described it as “the optimal revenue management of perishable assets through price segmentation.”

1.4 Origins of Revenue Management

The roots of modern revenue management can be traced back to the early days of the U.S. airline industry. Prior to the Airline Deregulation Act of 1979, fares for airline travel in the United States were regulated by the Civil Aeronautics Board (CAB). The CAB ensured that the airlines operated in a highly controlled environment designed to serve the public convenience and necessity.[9] The CAB required economic justification for any fares proposed by the airlines. Thus, there were few fares for customers to choose from. In the 1930’s all airlines offered all seats on a flight for the same price. But it was obvious to the airlines that passengers could be divided into two broad categories, based on their travel behavior and their sensitivity to prices. There were business travelers and leisure travelers. Business passengers tended to make their travel arrangements close to their departure date and stay at their destination for only a short time. There was little flexibility in their plans and they were willing to pay higher prices for tickets. Leisure travelers, on the other hand, booked their flights well in advance of their travel date. They stayed longer at their destinations and had much more flexibility in their plans. They would often decide not to travel rather than pay high fares. Since there was only one fare offered to both types of passengers, many of the leisure passengers chose not to fly, and many flights departed with empty seats.

Airline managers saw an opportunity to increase revenue by lowering fares in certain markets. The first experiment to offer low-fare service occurred in California on the San Francisco-Los Angeles route in 1940.[10] United Airlines began its Sky Coach Service using 10-passenger Boeing 247s and charging a one-way fare of $13.90. The CAB approved the low fares based on the lower operating cost of the B-247s and fewer amenities offered on board. The experiment was a success but ended shortly thereafter when the airline's fleet was turned over to the armed forces during World War II.

Throughout the next few decades, other discount fares were offered with varying degrees of success. First-class and coach-class became standard on all airlines. But the airlines were not permitted to offer different fares within the coach cabin and prices were set through a cost-plus pricing formula administered by the CAB. Carriers gradually became less efficient at operating their airlines, and coach fares rose over time as average costs increased.

During the 1960s, the CAB began approving new types of fares such as night coach fares and 7-21 day excursion fares based on length of stay. However, the airlines placed no limits on the number of seats that could be sold at these fares, and all were available on a first-come, first-served basis.

In the early 1970s, the CAB responded to demand for more discount fares by easing regulations for charter airlines. With their lower operating costs, the charter carriers were able to offer low fares and still earn a profit. For example, in the winter of 1976, passengers could travel from New York to Florida on a charter for as little as $99.[11] This fare was less than the average cost for a major airline to fly that market. So if the airline matched the charter fare, then it would lose money on the flight, even if it filled every seat.

This situation caused concern among the managers at the major airlines. Their initial thought was to figure out a way to reduce costs so they could remain competitive. But that was impractical. The costs of operating a major airline with its staffing and airport needs were simply much higher than the cost of running a charter operation. But then the executives at American Airlines realized something. On average, their planes were departing with half their seats empty. While the average cost of these seats was higher than the charter fares, the *marginal* cost was close to zero. So if they could find a way to sell just the empty seats at the charter fares, profits would increase dramatically. The challenge was to devise a plan that would make the empty seats available at the lower fair, while preventing passengers who were willing to pay the higher fare from buying low-fare seats. American Airlines' response to this challenge was the introduction of "Super Saver Fares" in 1977. With these fares came the beginning of modern day revenue management.[12]