Performance of Dynamic Programming Methods in Airline Revenue Management

by

Sarvee Diwan

B.Tech. (Civil Engineering)
Indian Institute of Technology Bombay (2006)

Submitted to the Department of Civil and Environmental Engineering
and Sloan School of Management
in partial fulfillment of the requirements for the degrees of

Master of Science in Transportation

and

Master of Science in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2010

Author . . . .
Department of Civil and Environmental Engineering
and Sloan School of Management
May 14, 2010

Certified by . . . .
Peter Paul Belobaba
Principal Research Scientist of Aeronautics and Astronautics
Thesis Supervisor

Certified by . . . .
Hamsa Balakrishnan
Assistant Professor of Aeronautics and Astronautics
Thesis Reader

Accepted by . . . .
Daniele Veneziano
Chairman, Departmental Committee for Graduate Students

Accepted by . . . .
Dimitris J. Bertsimas
Boeing Professor of Operations Research
Co-director, Operations Research Center

Performance of Dynamic Programming Methods in Airline Revenue Management

by

Sarvee Diwan

Submitted to the Department of Civil and Environmental Engineering
and Sloan School of Management
on May 14, 2010, in partial fulfillment of the
requirements for the degrees of
Master of Science in Transportation
and
Master of Science in Operations Research

Abstract

This thesis evaluates the performance of Dynamic Programming (DP) models as applied to airline Revenue Management (RM) compared to traditional Revenue Management models like EMSRb as DP models offer a theoretically attractive alternative to traditional RM models.

In the first part of this thesis, we develop a simplified simulator to evaluate the effects of changing demand variance on the performance of standard DP on a single flight leg. This simulator excludes the effects of forecast quality and competitive effects like passenger sell-up and inter-airline spill. In the next part of the thesis, we introduce two network based DP methods that incorporate the network displacement costs in the standard DP based optimizer and perform simulation experiments in a larger competitive network using the Passenger Origin Destination Simulator to study the performance of DP methods in airline Revenue Management systems.

The results of single flight leg experiments from the simplified simulator show that DP methods do not consistently outperform EMSRb and the sensitivity analyses show that the performance of DP relative to EMSRb depends on the demand variability, demand factor, fare ratios and passenger arrival pattern. The results from the PODS competitive network simulations show that DP methods, despite not showing any significant benefits in the simplified simulator, can outperform EMSRb when used in a competitive environment because DP's aggressive seat protection policy helps DP generate more revenues than EMSRb due to competitive feedback effects like inter-airline passenger spill-in, and passenger sell-up within the airline.

Thesis Supervisor: Peter Paul Belobaba
Title: Principal Research Scientist of Aeronautics and Astronautics

Acknowledgments

I would like to express my sincere gratitude to my research advisor, Dr. Peter Belobaba, for his guidance and inspirational presence throughout the course of my stay at MIT. This thesis would not be possible without his patience, continuous guidance and support. I really appreciate the research opportunity he has given me during my stay at MIT. His vast experience and knowledge about the Revenue Management in the airline industry made the work really enjoyable and highly enriching. I thank him for being such a wonderful advisor and mentor.

I must also thank Craig Hopperstad, the original developer of the Passenger Origin-Destination Simulator (PODS) on which my some of my results are based, for his technical support and patience in his explanations of sophisticated PODS techniques during our discussions. My understanding of the various intricacies in PODS would not be complete without his continuous support.

I extend my heartfelt gratitude to all the members of the PODS Consortium: Air France-KLM, Air New Zealand, Continental Airlines, LAN, Lufthansa-SwissAir, Delta-Northwest Airlines, SAS, United Airlines and Air Canada. The perspectives provided by various members during various PODS meetings were very helpful in shaping up the research work of this thesis.

I would also like to thank the lab mates at International Center for Air Transportation (ICAT) for being so helpful and making the stay at lab a thoroughly enjoyable experience. I would like to especially thank Claire, Olivier, and Chris for being such great travel companions during all PODS meetings. A special thanks to Claire for organizing numerous lab events which were always a welcome break from the routine stuff. A special thanks to Varun for being there in good as well as bad times over the last two years and for all the help he provided me with. Yannis, Nikolas, Thomas, Lishuai, Hanbong, and Alex were wonderful lab-mates who made the ICAT lab a fun place. I would also like to thank the administrators at CEE – Patty, Kris and Jeanette for being a great help throughout my stay.

My first year buddies, Harshad and Kshitij, made it really easy for me to settle down in this country when I first landed. A special thanks to Kshitij for those awesome Muddy sessions during our first year. I thank, Vibhu and Kushal, for being such amazing Edgerton-002 roomies and friends. I would cherish every single discussion we had during those late nights at our place. It is impossible to forget the help and support of Vikrant Vaze who has been there through thick and thin and has been a big source of support here at MIT. I will miss the amazing lunch and DuDo sessions we had over which we would discuss everything under the sun. I would also like to thank Vivek R., Ahmed, Abhishek, and Himanshu who have made the years spent at MIT a wonderful and memorable experience. This list would not be complete without my MST classmate and buddy Ta for making this journey a smooth ride.

Most importantly, I would like to thank my parents, without whose unwavering support, nothing could have been possible. Their unconditional love and support has helped me sail through difficult times. Finally, I thank god for giving me this opportunity.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The goal of this thesis is to evaluate the benefits of using Dynamic Programming based optimization models in revenue management for airlines. Airline revenue management has its roots in the post-deregulation era and this has gained widespread importance over the years. Revenue Management (RM) has presented industry and academia with new challenges since the advent of low-cost carriers because the assumptions made in the original systems no longer hold. The revenues generated by traditional RM models can be sub-optimal despite recent developments to improve them, and new optimizers that eliminate those assumptions may be necessary to attain the optimal revenues.

We will use Passenger Origin-Destination Simulator (PODS) as the simulation environment for our optimization models. It was developed by Hopperstad, Berge and Filipowski Hopperstad (2005) at the Boeing Company to model the airline booking process with competing airlines trying to maximize passenger revenues in different competitive network configurations. PODS has been further developed by the PODS Consortium, a research alliance between MIT and 9 leading international airlines.

1.1 Revenue Management in Airline Industry

In this section, we briefly introduce the concept of revenue management as applied to the airline industry. Revenue Management (RM) serves to design and manage service products to to maximize revenue (Weatherford, 1991). RM addresses basic questions like which segmentation or differentiation scheme to use, how to set prices across various categories of the product, and how to allocate capacity to different products or categories to maximize revenue (Talluri and Ryzin, 2004). In this thesis, we will use the term revenue management for just the capacity allocation aspect of this system. In the context of the airline industry, airlines strive to understand passenger's willingness-to-pay (WTP) so that they can maximize their revenues by selling their seats to high fare passengers while attracting sufficient low fare passengers so as to fill up the airplane. Hence, in terms of the airline industry, we define revenue management (or yield management) as the process of allocating seats to different fare categories on an airplane by making use of the knowledge of demand segmentation and customer's willingness-to-pay so as to maximize revenue.

Round Trip Fare	Class	Adv. Purchase	Min. Stay	Change Fee	Comments
458	N	21 Days	Sat. Night	Yes	Tue/Wed/Sat
707	M	21 Days	Sat. Night	Yes	Tue/Wed
760	M	21 Days	Sat. Night	Yes	Thu-Mon
927	H	14 Days	Sat. Night	Yes	Tue/Wed
1001	H	14 Days	Sat. Night	Yes	Thu-Mon
2083	B	3 Days	None	No	2 X OW
2262	Y	None	None	No	2 X OW
2783	F	None	None	No	First Class

Table 1.1: Example of BOS-SEA Restricted Fare Structure (Belobaba et al., 2009)

Today, airlines offer various fare products with different restrictions at different fare levels to cater to different segments of demand while maximizing revenues. At one level, they differentiate by laying out different cabins (economy, business class and first class cabins) and associating different levels of service for different cabins. At another level, they create differentiated products within each class by adding restrictions[1]. Before deregulation, airlines did not have much flexibility in deciding the fare structures and hence, competition was based on better quality of passenger service and providing more frequent flights. But soon after deregulation, airlines were

[1]Belobaba (1995) discusses *differential pricing* and its use in Air Travel Demand Segmentation.

free to decide their own fare structures and this led to the design of the fare structures with restrictions to segment passenger demand. The most common restrictions to segment the demand are: (1) Advance Purchase, (2) Saturday Night Minimum Stay, (3) Change Fee and (4) Limited Refundability on Canceled Bookings. An example of a restricted fare structure is shown in Table 1.1.

Differential Pricing is the practice of offering different fare products at different prices with varying characteristics for the same O-D market. *Revenue Management (RM)*, on the other hand, helps the airline decide how many seats to offer at a particular price for a flight. Typically, airlines use RM systems that set booking limits on low-fare seats by using the information available for schedules, demand and fare structures. Airlines realized that although the practice of differential pricing helps them gain revenues, that alone could not produce optimal revenues and hence there was a need for an RM system. In essence, the main objective of RM systems is to protect seats for late arriving high-fare passengers so as to strike the right balance between load factors[2] and yields[3] resulting in revenue maximization.

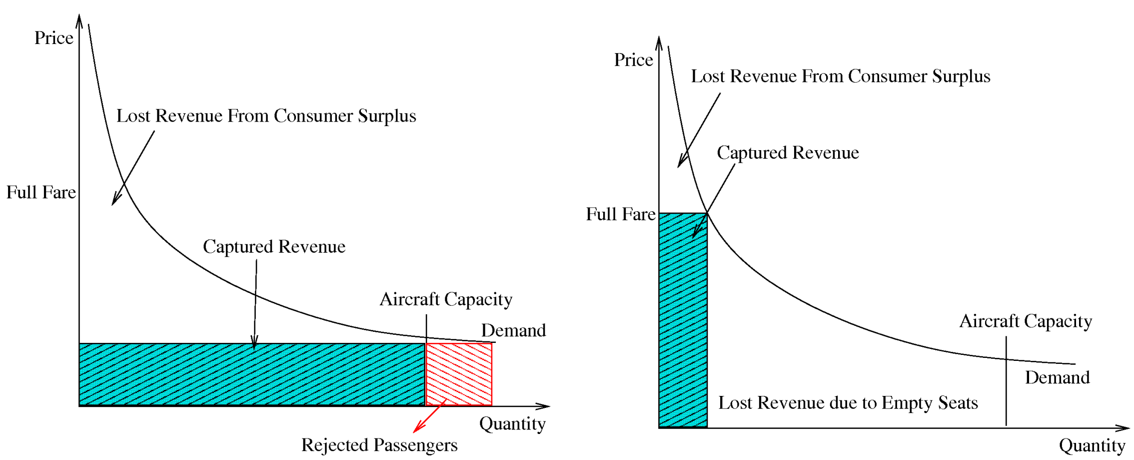

(a) Under-protection

(b) Over-protection

Figure 1-1: Revenue Implications for (a) Under-protection and (b) Over-protection

Traditionally, RM systems use mathematical models and computer infrastructure to address three basic problems in seat inventory control: (1) Overbooking, (2) Fare-

[2]Load Factor is defined as the ratio of Revenue Passenger Miles to Available Seat Miles
[3]Yield is defined as Revenue per Revenue Passenger Mile (Revenue per RPM)

Class Mix Control, and (3) Origin-Destination Control. Readers can refer to Barnhart et al. (2003) for an overview of applications of Operations Research (OR) in the airline industry. An airline seat is a perishable item - any unsold seat would represent a loss of revenue for the airline once the flight departs. Hence, too many seats saved for high-fare passengers could result in low load factors and loss in revenues, but on the other hand, booking too many low-fare passengers might result in some high-fare passengers being denied a seat and result in loss in revenues due to inability to capture consumer surplus. Figures 1-1(a) and 1-1(b) show the effects of dilution and over-protection of seats respectively.

Most airlines use a nested booking control RM system. Nested booking means seats are jointly protected for higher classes from all the ones below them. For example, if 10 seats are protected for fare class 1, then only those passengers who request this fare class can book these 10 seats. However, if 20 seats are protected for Class 3, then all passengers requesting either class 1, 2 or 3 can book those seats. It is easy to see that bookings for class 1 (highest fare) will always be accepted as long as there are unsold seats on the flight. Figure 1-2 shows the joint protection of seats in a flight with remaining capacity of 100 seats. We will discuss the evolution of RM

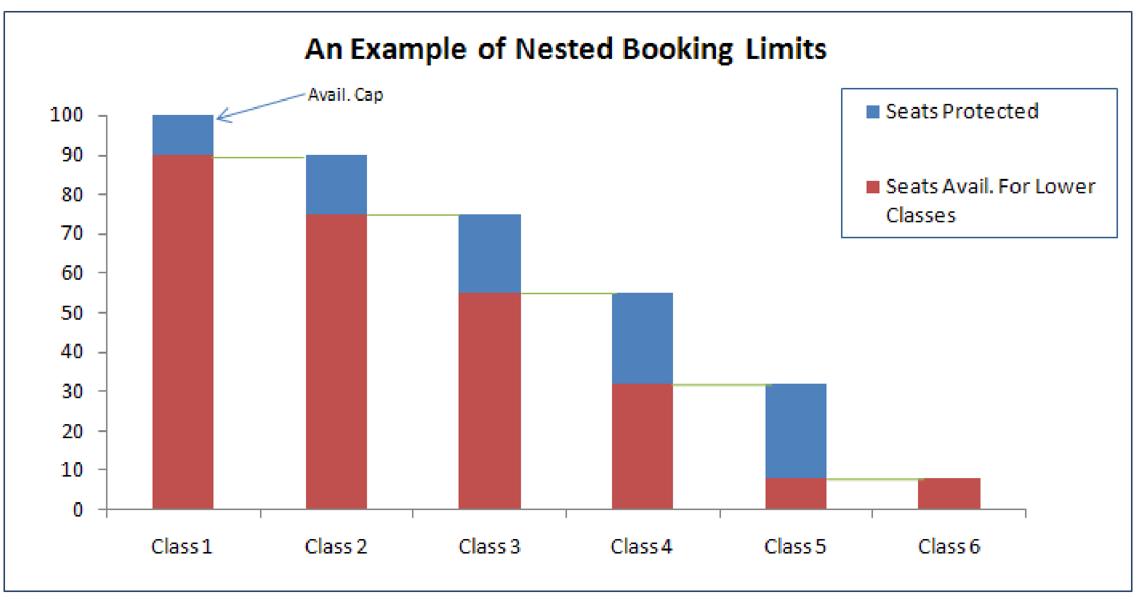

Figure 1-2: Nested Protection of Seats

systems and various models addressing the problem of seat inventory control in detail in Chapter 2.

1.2 Recent Developments in the Airline Industry

The deregulation of the airline industry served as a catalyst for the growth of low-cost carriers (LCC). Before deregulation, domestic US markets were served by a few select airlines because of high barriers to entry due to government regulations. There was no flexibility in deciding fares. Post deregulation, many new competitors entered the market and the Network Legacy Carriers (NLC), facing threat from the entry of new competitors, started using RM systems to increase their revenues and maintain market shares. Since then, the LCCs, with their simplified and less restricted fare structures, have grown very rapidly over the past few years and have captured a substantial portion of the domestic traffic. According to MIT's Airline Data project website[4], LCCs now capture 27% of United States' domestic traffic as compared to little over 11% in 2000[5]. The website also shows that the LCCs have lower operating unit costs (Cost per Available Seat Mile). It allows them to offer low-fare products with more frequency on popular routes. An example of a less restricted fare product offered by Southwest Airlines is shown in Table 1.2. It can be noted that the US

Round Trip Fare	Class	Adv. Purchase	Min. Stay	Non-Refund?	Comments
178	M	3 Days	1 Night	Yes	Special Sale
402	H	7 Days	1 Night	Yes	Tue/Wed
438	Q	14 Days	None	Yes	2 X OW
592	B	7 Days	1 Night	Yes	Off Peak
592	B	7 Days	1 Night	Yes	Peak
634	Y	None	None	No	Unrestricted

Table 1.2: Example of PVD-SEA Simplified Fare Structure (Belobaba et al., 2009)

domestic air traffic as well as passengers' WTP has been on a decline in this decade due to the 9/11 attacks, economic downturn in 2000 and the recent financial crisis. Increasing unit operating costs due to soaring labor and fuel costs, a decline in the

[4]http://airlinedataproject.mit.edu
[5]The comparison is based on Revenue Passenger Miles (RPM)

number of business class bookings and a decline in the average fare has adversely affected the profitability of airlines. Most of the large US airlines have undergone bankruptcy protection and now are in the process of reducing their unit costs along with increasing revenues.

The use of Information Technology (IT) has made it easier for the LCCs to advertise their low-fare product offerings. Passengers can now get more complete information about fares, special offers etc for travel on a particular date in any OD market. The transparency provided by on-line distributors like Travelocity, Expedia, Orbitz etc. enables a passenger to make a more informed decision by comparing fares of all the airlines competing in that market.

In response to these developments, Network Legacy Carriers (NLC) have also moved towards the simplified fare structures regime. The legacy carriers try to offer the same fare structure offered by LCCs in the markets where LCC is a competitor to avoid losing market share. The legacy carriers can offer this simplified fare structure by removing some restrictions altogether or lowering the Advance Purchase (AP) requirements. In the next section, we will discuss the effects of simplified fare structures on revenues and the need for development of new RM methods to cater to this new environment.

1.3 Thesis Motivation and Objectives

The success of low cost carriers (LCC) has had a significant impact on the way the legacy carriers operate throughout the world. Legacy carriers had been banking on traditional revenue management methods which exploit the demand segmentation attained with various fare restrictions in fare structures. The degree of demand segmentation is very different in LCC fare structures when compared with that of a legacy carrier. Due to their low cost structure, LCCs can manage to be profitable despite a relatively unrestricted fare structure but the legacy carriers are forced to match their fares to avoid losing market share[6].

[6]Gorin (2000) provides a comparison between the business models of legacy carriers and LCCs

Passengers tend to buy the lowest open class independent of their WTP since there are very few restrictions under the new simplified fare structure. This makes it very difficult for an airline to differentiate between price-oriented demand and product-oriented demand[7]. Hence, under less differentiated fare structures, the assumptions of traditional RM methods are violated since business passengers are able to buy down to the lowest open class. This leads to difficulties for traditional RM methods to forecast demand accurately because of less high-fare bookings observed and this, in turn, leads to less protection for those classes. For more information on "spiral-down", the reader is referred to the work of Cléaz-Savoyen (2005).

Most traditional RM systems assume independent fare class demand and a pre-determined and sequential order of booking arrivals. It means that as the departure date comes closer, the percentage of bookings made by business passengers increase and the sell-up probability[8] of discount passengers increase. However, if any of the above assumptions are not satisfied, the revenues generated by traditional RM methods may not be optimal. Dynamic Programming (DP) based approaches model each arrival as a stochastic variable[9] and hence, one would expect DP optimization models to produce better solutions since they do not assume anything about arrival sequence and are dynamic in nature[10].

A number of studies have been done regarding the applicability of DP optimization models applied to airline revenue management and the studies have shown mixed results[11]. Given the attractiveness of DP models in the current environment as described above, the fact that DP models do not consistently outperform traditional RM methods is not entirely obvious. This thesis aims at answering the question of effectiveness of DP models as a viable alternative to the traditional RM methods. The goal is also to study the effect of various assumptions made in the DP models and their effect on the performance of these models.

[7] For more information on this, please refer to Boyd and Kallesen (2004)

[8] Sell-up Probability is the probability that a passenger is willing to buy a higher fare ticket for the same flight when he/she is denied the requested fare class

[9] A Detailed literature review on these methods can be found in section 2.3

[10] We will discuss this in Section 2.3

[11] The works of Tam (2008) and Vanhaverbeke (2006) have been discussed in 2.3

1.4 Thesis Organization

The remainder of this thesis is organized as follows: Chapter 2 presents a literature review of RM methods in general and DP based models in particular. We start off by describing the evolution of RM methods and then focus on Dynamic Programming based approaches that have been discussed in literature. We wrap up the chapter with a little review of the forecasting methodologies that have been developed in literature.

Chapter 3 presents an example of DP applied to a single leg. This chapter will help us build intuition about the results obtained by applying DP to a RM system and we can build on the results obtained from this small network. In this chapter, we develop our own simulation environment and present the results obtained through our simulator on a single-leg example.

Chapter 4 introduces Passenger Origin Destination Simulator (PODS) and various models that will be used for simulating DP in a competitive network environment using PODS. We introduce two new network based approaches for managing seat inventory using Dynamic Programming methodology.

Chapter 5 presents the results of the PODS simulations. We discuss and analyze the results under different fare structures to study the applicability of DP methodology under various fare structures.

Chapter 6 summarizes the results and draws conclusions from them. We also present a short note of possible future research directions.

Chapter 2

Literature Review

Revenue Management literature dates back to the 1960s with researchers focusing on increasing load factors and yields. Post deregulation, there is abundance of literature focusing on optimizing revenues as opposed to maximizing load factors or yields. We focus on the available literature in the field of Revenue Management (RM) in general and Dynamic Programming (DP) models in particular. The first section of this chapter discusses the evolution of traditional RM methods, the second discusses DP models available in literature and the last section gives a brief overview of the available literature in forecasting models.

2.1 Evolution of Revenue Management Systems

As we have previously stated, the goal of an efficient revenue management system is to efficiently manage inventories given a fixed capacity by taking advantage of demand segmentation. The first use of revenue management in the literature was through overbooking models. Their aim was to decrease the loss of revenue because of passengers not showing up on the day of departure of the flight. Overbooking refers to the process of selling more seats than available so that the above described no-show loss can be minimized. An additional cost that the airlines faced, however, were the denied boarding costs. Aggressive overbooking means that some passengers could not be accommodated on their respective flights and hence needed to be compensated for

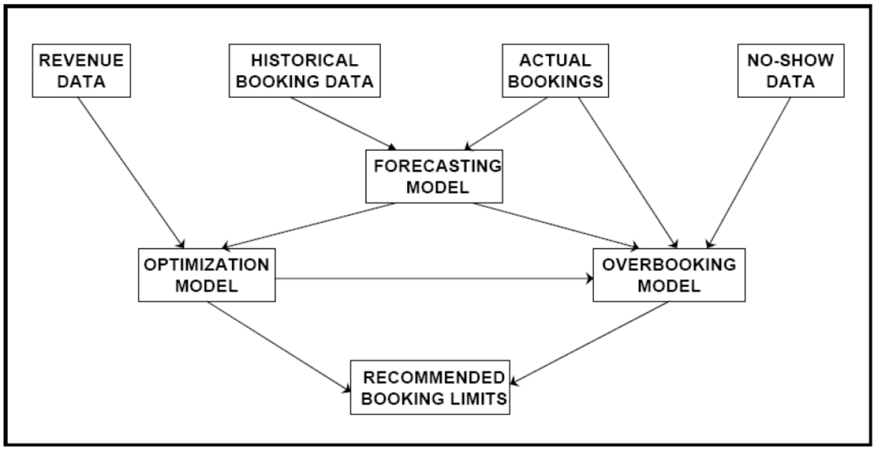

Figure 2-1: A Typical Third Generation RM System (Barnhart et al., 2003)

the same. For details on overbooking, readers can refer to Rothstein (1971, 1975).

Airline RM systems have evolved in terms of their mathematical modeling capabilities over the last thirty years. The first generation RM systems, developed in the early 1980s, were large database management systems in which airlines could store data obtained from observed demand and actual bookings. The decision making process was left to human analysts as these RM systems did not provide any guidance about the actions required to increase airline revenues. The next generation RM systems, developed in mid 1980s, allowed airlines to track actual flight bookings relative to an expected booking curve for the flight. These RM systems would issue an "exception reports" indicating a need for human intervention whenever a flight's actual booking deviated from the expected booking profile but provided no guidance about altering the booking limits to maximize revenues. In the late 1980s, the third generation RM systems were implemented by the advanced airlines. Apart from acting as a database like previous RM systems, these third generation RM systems helped the airlines in forecasting the demand and optimizing the seat allocation process for future flights. The major components of a typical third generation RM system are illustrated in Figure 2-1. The reader can refer to McGill and van Ryzin (1999) and Barnhart et al. (2003) for a thorough treatment of the evolution of RM systems.

2.2 Models for Seat Allocation

RM systems take the fare structure (i.e. different fare values for different classes alongside the associated restrictions) as input and try to optimally allocate seats for each fare class. The issue at hand is in managing the fixed and perishable inventory of seats so that a sufficient number of seats is saved for high-fare passengers while allocating the remaining seats at a discounted price to passengers who have a lower willingness-to-pay. There are two kinds of approaches to solve this problem in literature: Leg-based Controls and Origin-Destination Controls.

2.2.1 Leg Based Controls

Littlewood (1972) developed a model for the seat inventory control problem by considering a single-leg, two fare class problem. This model was extended to multiple legs and multiple fare types by (Buhr, 1982), (Richter, 1982) and Wang (1982). The approach common to all the above models is that of using serial nesting in fare classes and using the concept of expected seat revenue. Belobaba (1987a) provides a very comprehensive overview of the models discussed above.

Belobaba (1987b) and Belobaba (1989) developed a more general model that is applicable to any number of fare classes based on the concept of Expected Marginal Seat Revenue (EMSR). He subsequently developed a variant of this algorithm which allows joint protection of the upper classes from the next booking class right below them. This algorithm, EMSRb, has since become an industry standard for setting the booking limits on a flight leg basis. Optimal formulations for the same multi-class problem have been developed by Curry (1990), Brumelle and McGill (1993), Wollmer (1992) and Robinson (1995). Their results indicate that the booking limits generated by much simpler and computationally tractable EMSRb are very close to optimality. A detailed description of the EMSRb algorithm can be found in Belobaba (1992) and Belobaba and Weatherford (1996).

All the methods described above are leg-based control algorithms since they assume that all the legs carry one itinerary. However, in reality, airlines sell multi-leg itineraries and hence resources (seats on an airplane) are shared by different fare classes as well as between local and connecting passengers. RM methods which take this network effect into account are classified as Origin Destination controls.

2.2.2 Origin Destination (Network) Controls

The importance of O-D seat inventory control stems from the fact that maximizing of revenues on all individual flight legs might not mean that the total revenues for the entire network are maximized (Williamson, 1988). Since most of the large airlines in the world operate with a hub and spoke network, this type of control gains much more importance.[1]

There are two basic approaches to solve the problem of network O-D control. The first one is based on the idea of virtual nesting. In this approach, the idea is to group local and connecting fares in "virtual buckets" which are visible only to the airline and then apply leg-based booking limit control such as EMSRb to determine the booking limit of each of the virtual buckets using path forecasts (A path is defined as a collection of flight legs that form an itinerary between an origin and a destination within a network). The results showed that this initial approach of bucketing gave preference to high-fare connecting passengers despite the fact that one connecting passenger might be displacing two local passengers whose contribution to total network revenues might be higher than that of one connecting passenger. The reader is advised to refer to Williamson (1992) and Vinod (1995) for further details on this approach.

Wysong (1988) and Smith and Penn (1988) developed the Displacement Adjusted Virtual Nesting (DAVN) methodology in which the revenue input to the virtual buckets was adjusted for the network displacement costs incurred by a connecting passenger since it displaces a local passenger on a particular flight leg. The displacement costs are obtained by solving linear programming models (Williamson, 1992). After incorporating the network displacement costs, the problem can be solved at the leg-

[1] Belobaba (2002) provides examples of the shortcomings of leg-based control mechanisms in a network setting and explains how O-D control can be used to overcome them.

level using virtual buckets and EMSRb. A good overview of these models and their evolution can be found in Belobaba (2002).

The second approach, apart from virtual nesting, to solve the problems of O-D control is the use of bidprices as developed by Simpson (1989), Williamson (1992) and Smith and Penn (1988). The bidprice is the sum of marginal revenue values for an incremental seat on all legs of a given itinerary. Bid-Price Control provides a threshold number for an airline to accept a booking request — the request will be accepted if the bid-price is lower than the O-D fare requested, and will be rejected otherwise. Various bidprice algorithms are discussed in literature — Network Bid Price (NetBP) method, Heuristic Bid Price (HBP) method, and Prorated Bid Price (ProBP) method. A more detailed description of these bidprice algorithms can be found in Belobaba (2002) and Bratu (1998).

2.3 Dynamic Programming Models

Most traditional RM methods assume independent demand for various fare classes along with a predetermined and sequential arrival of demand for fare class bookings, typically lowest classes booking earlier than higher classes in the booking process. These methods, being static RM methods, set booking limits based on the total remaining demand with no distinction among future time frame demand levels. The booking limits do not change until the next re-optimization point in the booking horizon. Dynamic Programming (DP) based approaches, although assuming independent demands, do not assume any passenger arrival process and use the forecasts by time frame and by class to allocate seats. DP models are not static within a time frame of the booking process as they determine the booking limits as a function of bookings in hand and remaining capacity even within a time frame. The booking limit matrix, which is a function of number of future time frames and bookings in hand in each of them, helps the DP models adjust dynamically to changing passenger booking patterns even before the next re-optimization point occurs. DP models offer a theoretically attractive alternative to traditional RM models because they relax two out of the three assumptions made in traditional RM models.

Lee and Hersh (1993) develop a discrete-time DP model for finding an optimal booking policy. They allow for non-stationary demand which is modeled as a Poisson demand at the fare-class level. They model the arrival process as a Markov decision process where they divide the reservation process into multiple decision periods, each small enough that the probability of two or more passengers booking is zero, and use DP formulations to decide whether or not to accept the booking request (they allow group bookings). They also talk about the computational intractability of such an approach for a real-world application.

Gallego and Ryzin (1994) and Gallego and Ryzin (1997) develop a framework for dynamic pricing and control of inventories when the demand is stochastic and price sensitive. The demand is modeled as a stochastic process with intensity of demand as a function of the prices of products. Their approach of sub-dividing the booking horizon resembles that of Lee and Hersh (1993). The model calculates expected revenue based on the consideration of an accepted booking, a rejected booking and the possibility of sell-up. The objective of their model is to find the lowest open fare class at each time period. This greatly simplifies the control type that is stored in the RM system for each time-frame. For a more detailed description of this method (GVR-DP), refer to Vanhaverbeke (2006) and Tam (2008).

Bitran and Mondschein (1995) consider the problem of managing inventory when the customer arrival process is dynamic and stochastic and they propose a stochastic dynamic programming based model. They characterize the optimal policy as a function of capacity and time until the end of booking horizon. In their formulation, the stages correspond to the periods in which the planning horizon is divided and the state of the system at a time is characterized by the remaining capacity and the class of customer requesting a booking. They use heuristics to find upper bounds for optimal revenues so that the model remains computationally tractable.

Lautenbacher and Stidham Jr. (1999) develop a discrete time finite horizon Markov decision process (MDP) formulation to solve the single-leg inventory control problem. They use DP with a backward recursion to solve the model. Similar to previous approaches, they divide the booking horizon into sub-intervals so that no more than one request may arrive during one period.The decision to accept or reject a request is made depending on the remaining capacity, the remaining time in the reservation period and the requested fare class. Subramanian et al. (1999) extend this framework to include overbooking, cancellation and no-shows. The standard Lautenbacher DP method (LDP) incorporates all the common components of the previously discussed DP models and hence will be used as the fundamental DP formulation when we test the performance of DP models with traditional models in this thesis. LDP will be discussed in more detail in Chapter 3 when we discuss the models that we use in our simulations.

Zhao and Zheng (2000) propose a model for dynamic allocation of a perishable product over a finite time-horizon. They allow for the demand to be non-homogeneous. Zhao and Zheng (2001) extend the model and discuss its application to an airline inventory control system. They consider a two-class dynamic seat allocation model and derive the structural properties of th optimal policy. They construct the optimal policy as a threshold policy and compare it to an extended version of Littlewood's rule. They establish that their model based on Dynamic Programming produces policies similar to the ones obtained by applying Littlewood's formula. They show that although the policies generated by Littlewood's rule and its variants are not optimal, they are very close to optimal as long as the sell-up does not decrease with time.

Feng and Xiao (2001) consider a stochastic control model and develop optimal control rules based on the assumption that the demand follows Poisson distribution. They develop the model to be used in a hub-and-spoke environment. They discuss the idea of *Expected Marginal Time Cost* along with the idea of Expected Marginal Seat Revenue and they link the two together using the equation for calculating the bid price through a value function (similar to all previous papers).

Bertsimas and Popescu (2003) design dynamic optimization techniques based on stochastic demand for multiple classes using bid-prices from a linear programming relaxation. They argue that their algorithm performs better than the additive bidprice algorithms that are common in practice. They then extend their algorithm to incorporate no-shows and cancellations. They compare their algorithm, which approximates the opportunity cost by a piecewise linear approximation, with the traditional bid price control. They show that their policy works the same as a bid price policy when the load factors are low and better than bid price control when the load factors are high.

Bertsimas and de Boer (2005) combine a stochastic gradient algorithm and approximate dynamic programming to improve the booking limits obtained by any static nested booking policy. They show that considerable revenue gains can be obtained in a general network setting if the starting point fed to their model is the EMSR booking limit.

The PODS Consortium at MIT has been actively studying the problem of using DP models for revenue management in competitive scenarios (Vanhaverbeke (2006) and Tam (2008)). We will use PODS to simulate DP in a competitive network environment and will present results in Chapter 5.

Vanhaverbeke (2006) found that the DP formulation proposed by Lautenbacher and Stidham Jr. (1999) (LDP) generates only slightly better revenues than EMSRb in an unrestricted and simplified fare structure environment because it still assumes independent demand of fare classes. He also simulated revenue impacts by using the DP model as proposed by Gallego and Ryzin (1997) (GVR-DP) and found that it appears to be more promising but generates even worse results than other RM methods under competitive scenarios in simulations. He concludes that passengers willingness to sell-up is difficult to estimate because it depends on competitors fares and seat availability in future decision periods which vary with competition. He uses predetermined estimates of sell-up that do not adapt to competition against more advanced RM methods and thus lead to poor performance. He concludes that GVR-DP is very sensitive to forecasts of probabilities of sell-up.

Tam (2008) incorporates improved estimation of probabilities of sell-up in his simulation and concludes that in a single-leg network, LDP performs much better when fed with the right sell-up probabilities whereas GVR-DP performs inferior to other RM methods independent of the sell-up probabilities being used. However, he finds that in a much larger symmetric network with two competitors GVR-DP performs well as compared to other RM methods when accurate sell-up rates are estimated while LDP performs poorly. He concludes with the remark that the success of using theoretically appealing DP methods depend on the network structure, sell-up estimates and the competitor's RM systems.

We make an attempt to characterize the main models developed for the RM literature in Table 2.1 based on the two criteria enumerated below.

1. Policy: The policy discussed be either static or dynamic

2. Resources: The model could be developed for a single resource or for multiple resources (either a single-leg or a hub and spoke network)

	Static	**Dynamic**
Single Leg	Belobaba (1987b), Belobaba (1989), Wollmer (1992), Brumelle and McGill (1993), Robinson (1995),	Lee and Hersh (1993), Gallego and Ryzin (1994), Lautenbacher and Stidham Jr. (1999), Subramanian et al. (1999), Zhao and Zheng (2000),Zhao and Zheng (2001)
Network	Curry (1990)	Bitran and Mondschein (1995),Gallego and Ryzin (1997),Feng and Xiao (2001), Bertsimas and Popescu (2003), Bertsimas and de Boer (2005)

Table 2.1: Classification of RM Literature

2.4 Demand Forecasting Models

The demand forecast plays a central role in Revenue Management. The whole idea of revenue management is to allocate a resource with fixed capacity to the corresponding demand. We know the capacity perfectly well but need to forecast the demand. The performance of any RM method depends very heavily on the quality of forecast available to it so much so that even a very good optimizer when fed with a poor forecast will not be able to produce optimal revenues. Hence, we discuss some of the most important and frequently used forecasters in this section.

The forecast is usually based on some historical database consisting of bookings of previous flights and current bookings for future flights. It is very much possible that the observed bookings are less than the actual demand. It can happen when a particular fare class was closed and passengers were rejected for that fare class. Then the observed bookings/boardings will not contain the data of "true demand" and it remains unobserved. Hence, it is imperative that the booking data is detruncated to estimate the unconstrained demand and this demand should, then, be used to forecast future demand for that fare class.

2.4.1 Standard RM Forecasting

Pickup Forecasting is a more elaborate form of time-series forecasting. A time-series forecast would simply be a mean (weighted or unweighted) of final departure bookings on a set of similar flights. Pickup forecasting averages not just the final bookings but also incorporates more information on the number of passengers booked in the intervals before departure. It uses the average number of previous unconstrained bookings and the changes in bookings over time - i.e., the number of passengers that are "picked up" from time period to time period. The demand forecast can be obtained by adding the average pick-up from one time period to another for a fixed number of previous flights to the current bookings(or forecast). A detailed description of this methodology along with a numerical example can be found in Gorin (2000).

A slight variant to the traditional pickup forecasting is *Pickup Forecasting with exponential smoothing*. It allows putting more weights on recent samples with respect to the older ones.

Regression Forecasting is another classical forecasting method. This forecasting method is 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. Zickus (1998) argues that although pickup forecasting is a special case of regression forecasting, the regression forecaster tends to be slightly more accurate since total bookings on day of departure is correlated with total bookings some day before departure. For a detailed explanation on regression forecasting, refer to Zickus (1998).

2.4.2 Recent Developments to Address Spiral Down

The advent of simplified fare structure has allowed passengers to buy down to the lowest open class independent of their WTP since there are very few restrictions. This leads to difficulties for traditional RM methods to forecast demand accurately because of less high-fare bookings observed and this, in turn, leads to less protection for those classes. Cléaz-Savoyen (2005) discusses the phenomenon of "spiral down" in detail. The traditional RM forecasting methods discussed above would fail to provide accurate forecasts of demand by willingness to pay under simplified fare structure regime. There has been widespread interest and research on forecasting passenger demand for unrestricted fare structures. We present two of the most recently developed models in this section.

Hybrid Forecasting is based on two underlying concepts: (1) The notion of "price-oriented and product-oriented" demand as developed by Boyd and Kallesen (2004), and (2) 'Q' Forecasting for price oriented demand as developed by Belobaba and Hopperstad (2004). The philosophy behind the methodology is that there are basically two kinds of demand - product-oriented demand, which is the demand for a certain fare-class because consumers like it for its attributes and the price-oriented demand, which is the demand for the cheapest available product. Hybrid forecasting forecasts these two demands separately and then combines them to estimate total forecasts.

Q-forecasting relaxes the assumption of independent demands for different fare-classes by incorporating probability of sell-up, i.e., the probability that a passenger is willing to buy a higher fare ticket for the same flight when the passenger is rejected for the demanded fare class. The challenge here is to estimate the sell-up probabilities to obtain a Q-equivalent number of bookings. Cléaz-Savoyen (2005) discusses various methods that can be used to estimate sell-up probabilities. Combining this concept with that of demand segmentation as price-oriented or product-oriented, we can use regular pick-up forecasting to estimate the product-oriented demand and then use Q-forecasting to estimate price-oriented demand. Detailed descriptions of the methodology can be found in Belobaba and Hopperstad (2004),Cléaz-Savoyen

(2005) and Reyes (2005).

2.5 Chapter Summary

We motivated the use of DP as an optimizer in RM systems by comparing the assumptions of traditional RM methods with those in DP models and reviewed some DP models published in the literature. We also reported that DP, despite being the theoretical model of choice, did not always outperform traditional RM methods in some previous studies[2]. In the next chapter, we develop our own simplified simulation environment to study the performance of DP in a single leg network with no competition.

[2] The works of Tam (2008) and Vanhaverbeke (2006) were discussed

Chapter 3

Simulation of DP and EMSR in a Single-leg

We represent the problem of single-leg seat inventory control as a Dynamic Programming (DP) problem along the same lines as proposed by Lautenbacher and Stidham. We describe the model in this chapter and then develop a simulation environment to test the performance of DP in a single-leg case by changing the variability in demand. The revenues obtained by using DP are contrasted with those obtained by using traditional RM methods (EMSRb). This highly simplified simulation experiment gives us some fundamental insights into the performance of DP and will be used as a basis to explore the performance of DP in a more comprehensive simulation environment such as PODS in the following chapters.

3.1 Simulation Environment

Walczak (2006) reports that practitioners are sometimes reluctant to use RM methods based on DP since it is difficult to model high-variance, real-life demand as most of the existing DP models assume that the passengers in each fare class arrive according to an independent Poisson process. This assumption forces the demand variance to be equal to the mean demand which can be a serious under-estimation of the variance (or equivalently, a huge over-estimation of the forecast quality). We develop a simulator to test the effect of Poisson assumption in DP models by varying z-factors[1] (or demand variability) while keeping the mean demand constant for each fare class. The simplicity of the single-leg case provides us with more insights on the performance of DP as compared to EMSRb when both of them are tested under the same scenarios. An evaluation of this kind is valuable since it eliminates the effect of forecast quality, competition and network effects and hence, we believe that it can serve as the purest form of evaluation of DP as an optimizer in RM systems.

We divide the simulation process into four main stages - Passenger Generation, Forecasting, Optimization and Passenger Booking. We discuss these four stages in the sections that follow.

3.2 Passenger Generation Process

A suitable model for modeling the passenger arrivals is necessary as a first step before we can proceed with the task of evaluating various optimization techniques for airline RM systems. The model should allow us to generate a random number of passengers at different times of the booking process and hence, it should allow for the fact that passenger arrival intensities may vary over time. According to the overview given by McGill and van Ryzin (1999), the Poisson process is, by far, the most popular dynamic model of airline demand but it has a serious drawback – the mean, and the variance in a Poisson distribution are equal, which need not hold in practice. There are numerous models available in literature for modeling customer arrival process such as Bertsimas and de Boer (2005) and Weatherford et al. (1993). We model the arrival process along the same lines as Bertsimas and de Boer (2005).

To model the arrival process in our simulation keeping above factors in mind, we use a Non-Homogeneous Poisson process (NHPP) with arrival intensity given by:

$$\lambda_f(t) = A_f \cdot \beta_f(t)$$

[1] We define z-factor as the ratio of the variance to the mean. $zFactor = \frac{\sigma^2}{\mu}$

$$\beta_f(t) = \frac{1}{\tau} \cdot \left(\frac{t}{\tau}\right)^{\alpha-1} \cdot \left(1 - \frac{t}{\tau}\right)^{\beta-1} \cdot \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}$$

where

$$A_f \sim Gamma(\gamma, \delta)$$

$$\beta_f(t) \sim Standardized\ Beta(\alpha, \beta)$$

A good characteristic of this model is that the mean and variance can be decoupled from each other and it would help us simulate the effects of varying variances on the performance of optimizers being tested while keeping the mean demand at the same level. Intuitively, A_f determines the mean expected demand for fare class f and the arrival pattern is modeled by the standardized beta distribution β_f which will be useful in distributing various demand levels across time-frames before departure. A plot of beta distribution for different values of α and β is shown in Figure 3-1.

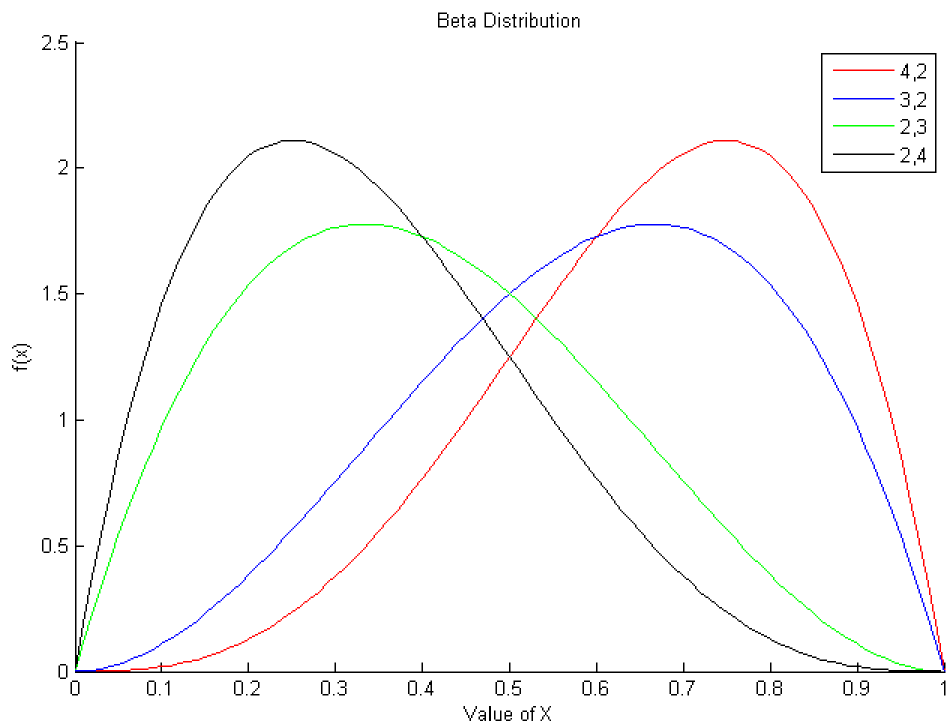

Figure 3-1: Plot of Beta Distributions for Different (α, β)

It can be seen that by varying α and β, we can vary the shape of beta distribution which, in turn, can be used to represent the arrival patterns of different fare classes.

For example, it has been empirically observed that low-fare leisure passengers tend to arrive early in the booking process and high-fare business passengers arrive late in the booking process (partly because of the way they plan their trips and partly due to the advance purchase restrictions in fare structures). The standardized beta distribution allows us the flexibility to model these practical aspects of passenger arrival process.

To define the model, we need to to specify four variables for each fare class:

- γ, δ: Required to specify the Gamma distribution and they determine the mean expected demand for each fare class

- α, β: Required to specify the standardized Beta distribution and they define the shape of the arrival curve (by fare class)

3.2.1 Statistical Properties of the Total Demand

Based on a monograph by Grandell (1997), a discrete random variable, N, is said to be *mixed Poisson distributed*, with *structure distribution* U if

$$p_x = E\left[\frac{\Lambda^x}{k!}e^{-\Lambda}\right] = \int_0^\infty \frac{\lambda^x}{x!} \cdot e^{-\lambda} \cdot dU(\lambda)$$

In the model used above, we have used the structure distribution as a Gamma distribution which gives us

$$U(\lambda) = \frac{\lambda^{\gamma-1}}{\Gamma(\gamma) \cdot \delta^\gamma} \cdot e^{-\lambda/\delta}$$

where $\Gamma(\gamma)$ is the gamma-function, defined by

$$\Gamma(\lambda) = \int_0^\infty x^{\gamma-1} e^{-x} dx, \quad \gamma > 0$$

γ is called the shape parameter and δ the scale parameter. Hence, we can say that Λ is $\Gamma(\gamma, \delta)$. Replacing $dU(\lambda)$ using above equation and rearranging terms, we get

$$p_x = \int_0^\infty \frac{e^{-\lambda}\lambda^x}{x!} \cdot \frac{e^{-\lambda/\delta}\lambda^{\gamma-1}}{\Gamma(\gamma)\delta^\gamma} d\lambda$$

Integrating by parts, we get

$$p_x = \frac{1}{\Gamma(\gamma)\delta^\gamma x!} \cdot \frac{\delta^{x+\gamma}}{(x+\gamma)(1+\delta)^{x+\gamma}} \cdot \int_0^\infty \frac{1+\delta}{\delta} e^{-\frac{1+\delta}{\delta}\lambda} (\frac{1+\delta}{\delta}\lambda)^{x+\gamma} d\lambda$$

Substitute $\frac{1+\delta}{\delta}\lambda$ by y, we get the following expression for p_x

$$p_x = \frac{1}{\Gamma(\gamma)\delta^\gamma x!} \cdot \frac{\delta^{x+\gamma}}{(x+\gamma)(1+\delta)^{x+\gamma}} \cdot \int_0^\infty e^{-y} y^{x+\gamma} dy$$

$$p_x = \frac{1}{\Gamma(\gamma)\delta^\gamma x!} \cdot \frac{\delta^{x+\gamma}}{(x+\gamma)(1+\delta)^{x+\gamma}} \cdot \Gamma(x+\gamma+1)$$

$$p_x = \frac{1}{\Gamma(\gamma)\delta^\gamma x!} \cdot \frac{\delta^{x+\gamma}}{(x+\gamma)(1+\delta)^{x+\gamma}} \cdot \Gamma(x+\gamma+1)$$

Using the property for Gamma functions that $\Gamma(y+1) = y\Gamma(y)$, we get

$$p_x = \frac{\Gamma(x+\gamma)}{\Gamma(\gamma)x!} \cdot \frac{\delta^x}{(1+\delta)^{x+\gamma}}$$

The above expression corresponds to the well-known negative binomial distribution $NB(r,p)$ with $r = \gamma$ and $p = \frac{\delta}{1+\delta}$. Hence, the total demand follows a negative binomial distribution with mean equal to $\gamma\delta$ and variance equal to $\gamma\delta(1+\delta)$. As can be seen from these expressions that the model used above for arrival process gives us the flexibility to model means and variances which are not necessarily the same (unlike Poisson process).

After studying the theoretical properties of the model for passenger arrival process, we turn our attention to simulating the model described above. We can not just use the Poisson random number generators with time varying intensities (based on the model above) because that would mean that within each time period we are forcing the means and the variances to be the same. Kuhl and Wilson (2009) provide a comprehensive overview of the literature available on simulating NHPPs. We used the thinning algorithm as proposed by Lewis and Shedler (1979).

3.2.2 Thinning Process

The thinning algorithm is one of the most commonly used methods for generating NHPP arrivals despite the concerns that it might not be very efficient computationally. There has been a lot of research in developing a more efficient accept-reject criterion. We have, however, implemented the textbook algorithm based on the following idea as presented in Kimms and M$\ddot{u}$ller-Bungart (2007): Assume that $\lambda(t) < \infty$ for all $t \in [0, T]$. Hence, we can find a value $\overline{\lambda}(t)$such that $\lambda(\tau) < \overline{\lambda}(t)$ for all $\tau \in [t, T]$. Now, suppose that the arrival times for the interval $[0, t]$ have already been generated, we simulate a homogeneous Poisson process with the intensity rate $\overline{\lambda}(t)$. It is obvious that this homogeneous process has got more arrivals than the NHPP arrivals, so thin out the stream of arrivals by rejecting an arrival at time τ with probability $\lambda(\tau) < \overline{\lambda}(t)$. The complete simulation technique can be described as follows: Given the values of $\overline{\lambda}(t)$, we generate a sequence of independent random variables $X_1, X_2, \cdots$ where X_i is a random variable with exponential distribution with rate $\overline{\lambda}(t)$ such that $t_i = \sum_{j=1}^{i-1} x^j$. We stop, as soon as we have reached a sequence of length N such that $\sum_{i=1}^{N} x^i > T$ and $\sum_{i=1}^{N-1} x^i < T$. The counting process having arrival times $\sum_{j=1}^{i} x^j, i \in I$ constitutes the desired NHPP with rate λ(t), where I is the set of indices such that

$$I = \left\{ i \in \{1, 2, \cdots, N-1\} : u^i \leq \frac{\lambda \sum_{j=1}^{i} x^j}{\overline{\lambda}(t_i)} \right\}$$

The thinning algorithm is more effective when $\overline{\lambda}(t_i)$ is as small as possible. In our implementation, we use $\overline{\lambda}(t) = max_{\tau \in [t;T]} \lambda(\tau)$. With the intensity function used, it is easy to calculate the maximum since we know that for $\alpha > 1$ and $\beta > 1$the beta distribution has a unique peak at $\frac{\alpha-1}{\alpha+\beta-2}T$.

We use the thinning algorithm with the intensity function described above to simulate the passenger arrival process as a Non-Homogeneous Poisson Process. We discuss the demand forecasting methodology in the next section.

3.2.3 Passenger Generation Results

We implement the thinning algorithm to simulate NHPP as described above. We look at a sample of results obtained from the passenger generation process and compare them to our apriori hypotheses about them. We will use the terms “passenger arrivals” and “passengers generated” in an interchangeable manner.

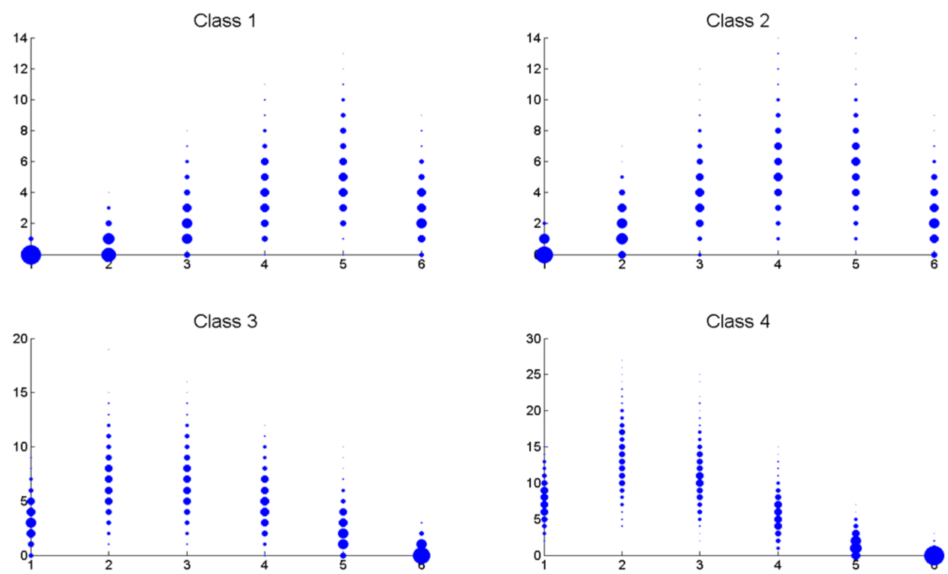

Figure 3-2: Passenger Arrival Samples Generated in 1000 runs for z Factor = 1

In Figure 3-2, the dot size represents the frequency of a particular demand level for a given fare class. We can make the following observations about the passenger arrivals:

- Almost no higher class (Classes 1 and 2) passengers arrive in the first time frame and no lower class arrivals (Classes 3 and 4) in the last time frame in almost all runs

- The simulated passenger distribution across time frames is similar to the input beta distribution (as shown in Figure 3-1)

Figure 3-3 shows the frequency distribution of total passengers generated. This plot experimentally verifies our claim that the total demand generated follows a negative binomial distribution with the specified mean.

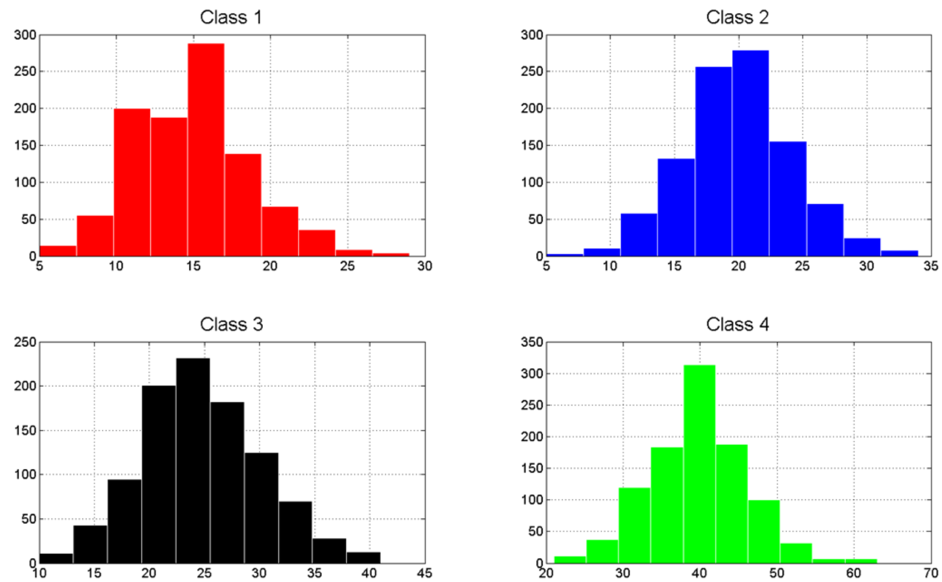

Figure 3-3: Frequency Distribution of Total Passenger Arrivals in 1000 runs for z Factor = 1

The total demand generated by time frame and by fare class is shown in Figure 3-4. It can be seen that most of the lower class passengers arrive in the initial time frames and most of the higher class passengers arrive late in the booking process.

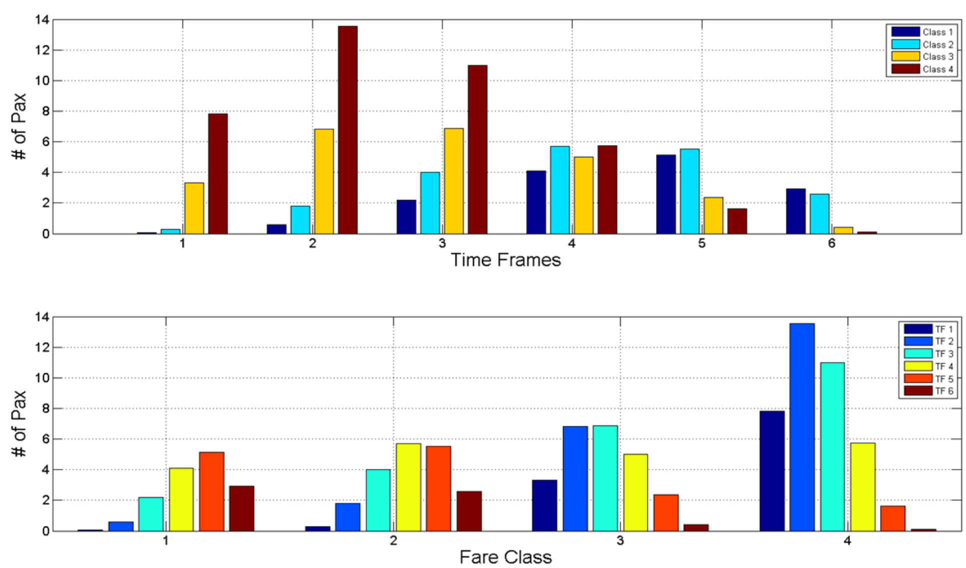

Figure 3-4: Total Passenger Arrivals by Fare Class and by Time Frame in 1000 runs for z Factor = 1

The results for simulation of the passenger arrival process show that we can simulate the real-life observation that most of the high-fare business passengers arrive late in the booking process and the low-fare leisure passengers book earlier in the booking horizon. We also observed that the thinning algorithm indeed generates NHPP in which the total demand generated is negative binomial distributed with the statistical properties as discussed in the preceding sections.

3.3 Forecasting

The basic idea of the simulator is to compare two optimizers (DP and EMSRb) and hence the forecasts need to be the same for both to have a fair comparison. The idea is to generate a large sample of passenger arrivals based on the specified distribution and treat that sample as a historical database of previous flights so that the expected simulated demands actually match the forecasts provided to the two optimizers.

We acknowledge that airlines use more sophisticated forecasting techniques as outlined in Section 2.4 but it is important for us to keep the forecasts same for both the optimizers to eliminate the effects of forecast quality on the optimization process. Since the basic idea here is to compare the optimization approaches, we use a very simple "forecaster" - For each run of the simulator, we generate a large number of samples beforehand and treat them as a historical database of flights for the past n years. We take their averages over time-frames and over fare classes and store them in a matrix which remains constant throughout the simulation run. This matrix is, then, fed as the mean forecast to both the optimizers. The benefit of using such an approach is that the effect of forecasting accuracy is eliminated and hence we are able to compare the optimizers on an "all-else-equal" basis.

Table3.1 represents a hypothetical forecast matrix for a case of 4 fare classes with 6 time frames before departure.

As we will see in the next section, DP requires only the mean forecast by time frame and by fare class. We feed this entire matrix into the DP and as time frame reduces, we reduce one column from left per time frame and feed the remaining matrix

Class/TF	1	2	3	4	5	6	Total
1	0.044	0.608	2.234	4.301	5.065	3.031	15.283
2	0.304	1.858	4.052	5.628	5.354	2.64	19.836
3	3.301	6.817	6.964	4.955	2.31	0.399	24.746
4	7.985	13.673	10.954	5.797	1.68	0.126	40.215

Table 3.1: Sample Historical Database for z=3.0

to DP. EMSRb requires the mean total forecast as well as the standard deviation of the demand, so we sum across all the existing columns to get the mean demand and the value of standard deviation is calculated as $\sigma = \sqrt{zFactor \cdot \mu}$.

We demonstrate this methodology to calculate the mean forecast, μ_f and the standard deviation, σ_f, for fare class f using the data from Table 3.1:

Time Frame 1:

DP: Entire Matrix is used as the mean forecast

EMSRb: $\mu_f = [15.2, 19.8, 24.7, 40.2]; \sigma_f = [6.7, 7.7, 8.6, 10.9]$

Time Frame 4:

DP: Columns 4, 5 and 6 along with all the rows are used as input

EMSRb: $\mu_f = [12.4, 13.6, 7.6, 7.6]; \sigma_f = [6.1, 6.4, 4.8, 4.7]$

3.4 Optimization

We compare two optimization models through our simulator - DP and EMSRb. As discussed in Section 2.2.1, EMSRb is the most popular model being used by airlines for managing their seat inventory on a flight leg. We evaluate the benefits/drawbacks of using DP as opposed to EMSRb for seat inventory control. Next, we discuss these two models in details.

3.4.1 Standard Lautenbacher DP

(Lautenbacher and Stidham Jr., 1999) proposed a discrete-time, finite-horizon Markov Decision Process (MDP) to solve the single-leg inventory problem. This model incorporates important components common to all existing DP models as described in previous chapter and hence, could be used to test the revenue impacts of using DP against traditional methods like EMSR. The standard LDP model does not take into account the effects of overbooking, cancellation, or no-shows.

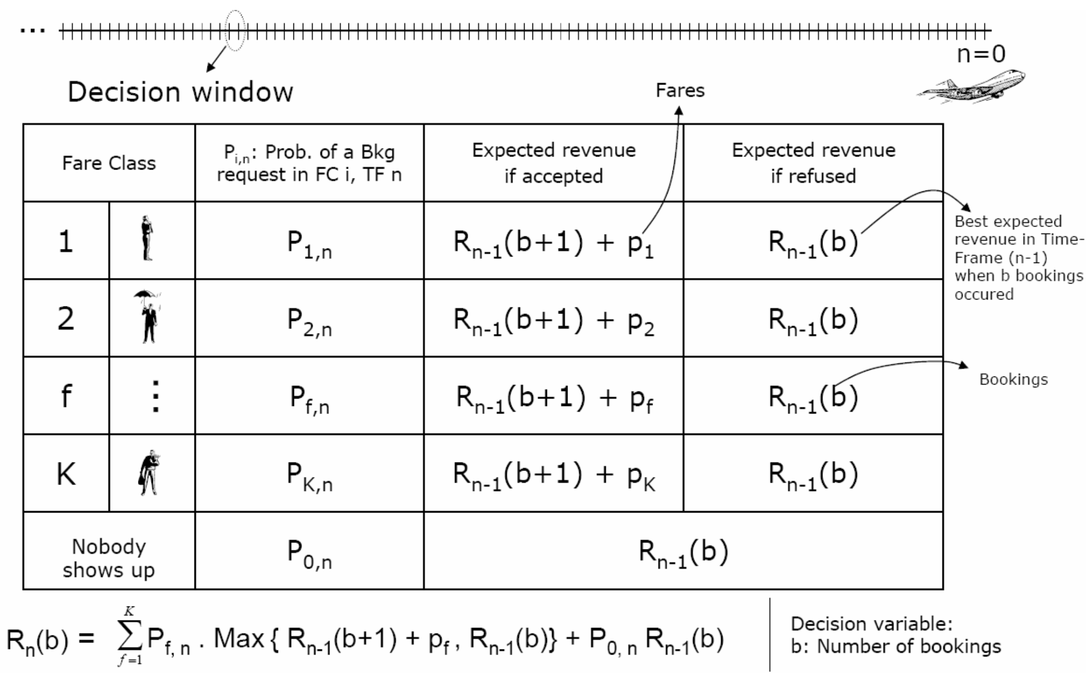

Figure 3-5: Standard Lautenbacher DP (Vanhaverbeke, 2006)

The booking period is divided into N decision periods so that the probability of two or more requests in any decision period is negligible. The decision periods are numbered in reverse order, with period N corresponding to the start of the booking period, and period 0 corresponding to the scheduled departure time as shown in the

Figure 3-5. We will use the notation as shown below.

C	Capacity of the Single Flight Leg
N	Number of Decision Periods
K	Number of Fare Classes
p_k	Fare of Class k
b	Number of Bookings Accepted
$R_n(b)$	Maximum Expected Revenue
$P_{f,n}$	Probability of a request for fare class f in decision period n
$P_{0,n}$	Probability of no request for any fare class in decision period n
$\Delta_n(b)$	Expected Marginal Seat Revenue of $(b+1)^{th}$ seat in decision period n
$B_{f,n}$	Optimal Booking Limit for fare class f in decision period n

Fare class demand is modeled as a Poisson process and each of the K fare classes, $1, 2, \ldots, K$ may arrive at anytime throughout the booking period. The decision to accept or reject any request depends on remaining capacity, remaining time to departure, and the revenue contribution of the requested fare class.

The maximum expected revenue for a given time frame with certain realized bookings is calculated by accounting for the probability that an accepted booking, a rejected booking, or no booking request may occur. The idea is that a request for fare class f will arrive with a probability $P_{f,n}$ during the decision period n and no booking request for any fare class occurs with a probability $P_{0,n}$. If the request is accepted, fare p_f is contributed to the expected revenue for the next decision period *n-1*. Hence, the maximum expected revenue for the next time frame is R_{n-1}(b+1) if the request is accepted, and R_{n-1}(b), if the booking request is rejected.

$$R_n(b) = \sum_{f=1}^{K} P_{f,n} \cdot max\{R_{n-1}(b+1) + p_f, R_{n-1}(b)\} + P_{0,n} \cdot R_{n-1}(b)$$

with the following boundary conditions

$$R_0(b) = \left\{ \begin{array}{cc} 0 & b \leq C \\ -R(b-C) & b > C \end{array} \right\}$$

where

$$R \geq max_f\{p_f\}$$

At the arrival of each booking request, the potential revenue generated with accepting the request is compared with the expected future revenue loss due to the removal of that seat from the available capacity. The expected marginal seat revenue of the $(b+1)^{th}$ seat in decision period $(n-1)$ when there are b realized bookings can be defined as follows:

$$\Delta_{n-1}(b) = R_{n-1}(b) - R_{n-1}(b+1)$$

$$B_{f,n} = min\{b > 0 : \Delta_{n-1}(b) > p_f\}$$

where

$$B_{f,n} \leq C$$

Optimal booking limits, $B_{f,n}$, for LDP are produced by backward recursion across the decision periods. The policy is to accept a class f request in decision period n if and only if $0 \leq b \leq B_{f,n}$.

3.4.2 Expected Marginal Seat Revenue - EMSR

Belobaba developed EMSRb, a seat inventory control model based on the concept of Expected Marginal Seat Revenue, as a revised version of his earlier model (Belobaba, 1992). This has since become an industry standard for managing seat inventory on a flight leg level. In his Ph.D. dissertation, Belobaba extended the two class model, proposed by Littlewood, to multiple fare classes with the Expected Marginal Seat Revenue (EMSR) heuristic (Belobaba, 1987b). This algorithm utilizes demand forecasts by fare class on a leg-basis to compute leg-based seat protection levels for nested booking classes. EMSR determines booking limits based upon the expected marginal revenue the probability of selling an additional seat in a given fare class multiplied by the revenue gained from selling that seat. As the number of seats protected in a particular fare class increases, the probability of selling that next seat decreases; thus, the booking limit for a fare class is determined when the EMSR is equal to the fare of the next lower class. As detailed below, EMSRb determines seat protection levels jointly for all higher classes relative to a given lower class, based on a combined total demand forecast and a weighted (based on total expected demand for each class) average fare for all classes above the class for which the booking limit is being calculated.

We define R_i as the fare for class i, $\overline{P_i}(S_i)$ as the probability that X_i, the demand for class i, is greater than S_i, the number of seats available in that class. The expected marginal revenue of making the S_i^{th} seat available is equal to the product of the above probability with the revenue for that class. Hence, we have the following

$$\overline{P_i}(S_i) = P(X_i \geq S_i)$$

$$EMSR_i(S_i) = R_i \cdot \overline{P_i}(S_i)$$

For a general class n, the following calculations are required to determine the booking limits.

$$\overline{X}_{1,n} = \sum_{i=1}^{n} \overline{X}_i$$

$$\hat{\sigma}_{1,n} = \sqrt{\sum_{i=1}^{n} \hat{\sigma}_i^2}$$

$$R_{1,n} = \frac{\sum_{i=1}^{n} R_i \cdot \overline{X}_i}{\overline{X}_{1,n}}$$

We need to find the value of π_n such that $EMSR_{1,n}(\pi_n) = R_{n+1}$ or $R_{1,n} \cdot \overline{P}_{1,n}(\pi_n) = R_{n+1}$. The booking limit for Class $n+1$ is determined as $BL_{n+1} = Capacity - \pi_n$.

A detailed discussion on the mechanics of EMSRb can be found in Belobaba (1992) and Belobaba and Weatherford (1996).

3.5 Passenger Booking Process

This part of the simulator acts like a "book-keeper". It takes booking limits obtained from the optimization process as input and books (accepts/rejects) passengers based on the booking limits. Alongside booking passengers, it keeps track of the statistics like revenues, load factors and the fare class mix.

3.6 Experimental Set-up

The main purpose of the simulation is to gain insights into the performance of DP and compare that with EMSRb.

	Fare (USD)	Mean Demand
Class 1	400	15
Class 2	300	20
Class 3	200	25
Class 4	100	40

Table 3.2: Fare Structure and Mean Demand for 4 fare classes

We simulate a single flight leg with one airline (and hence no competition). There are 4 fully differentiated fare classes and the booking horizon is divided into 6 time frames. The capacity of the single leg is 100 seats with mean demand across all classes summing up to 100. The fares and the mean demand of all four classes are as shown in Table 3.2. We simulate the above mean demand with varying *z-factor* for each class. It is important to note here that the *z-factor* of 1 corresponds to a Poisson distribution.

Figure 3-6 depicts the overall process flow for the simulation whereas Figure 3-7 represents the steps involved in simulating DP and EMSRb during the optimization process.

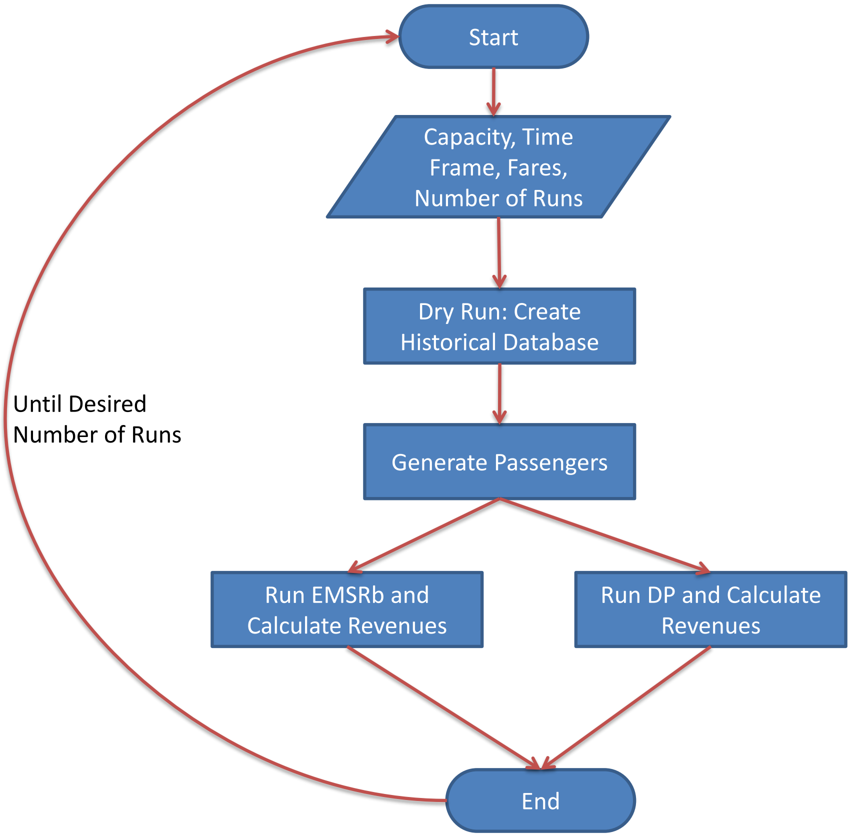

Figure 3-6: Overall Simulation Flowchart

3.7 Simulation Results

We present the results for the simulation experiment under the settings described in the previous section. We simulate two scenarios through our passenger arrivals simulation. The first case conforms to the practical reality in airline booking process where the leisure (or low fare) passengers arrive early in the booking process and the business (or high fare) passengers arrive in later time periods. In the next set of experiments, we relaxed this assumption so that any passenger is equally likely to arrive in any time period before the departure date.

3.7.1 Interspersed Arrivals - Lower Class Arrives Earlier

We require 4 parameters for each fare class to define its arrival pattern – two parameters for the gamma distribution (determines mean expected demand) and two for

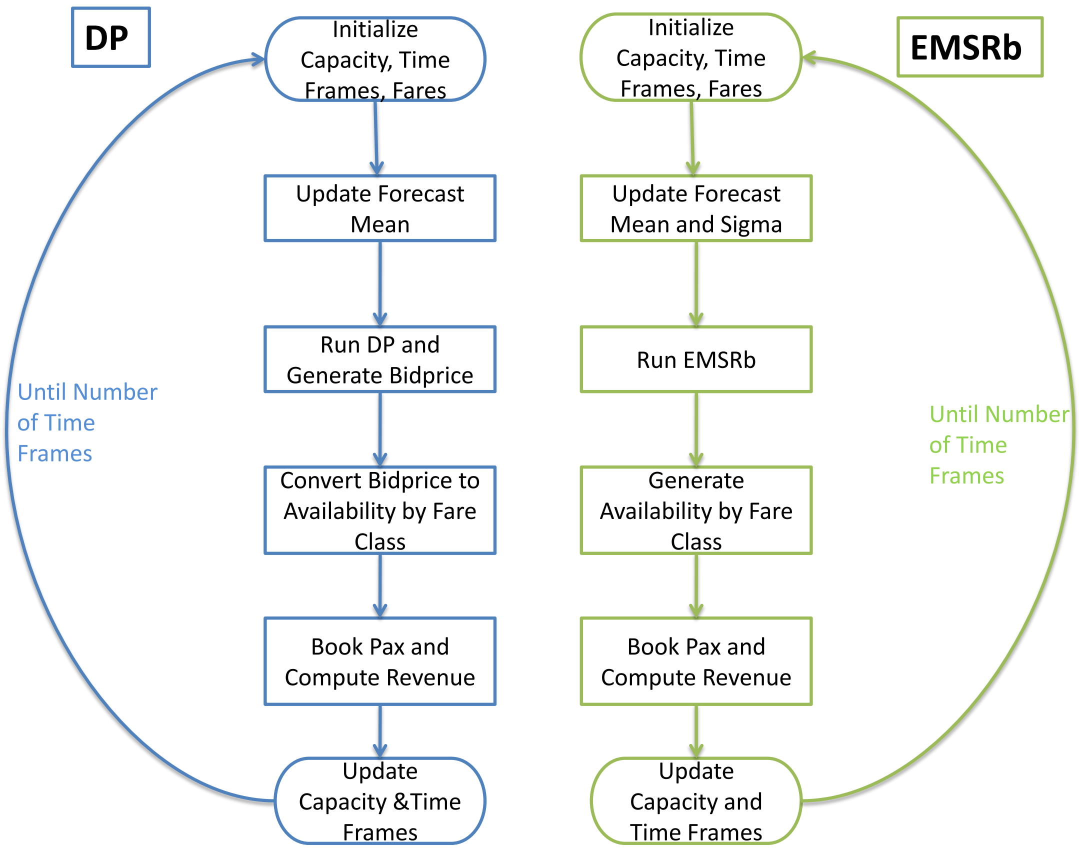

Figure 3-7: Overall Simulation Flowchart

the beta distribution (the shape determines the timing of arrivals during the booking process). The parameters for gamma distribution can be chosen very easily since we know the mean expected demand for each fare class as well as the variance we want to simulate. Hence, we get two equations (the mean and the variance as described earlier) in two unknowns which has a unique solution equal to $(\frac{\mu}{zFactor-1}, zFactor-1)$. Once we fix the parameters for gamma distribution, we only need to specify the parameters for beta distribution. The ideal way to determine these parameters would be to calibrate these so that we get a booking curve which matches what we observe in reality. The main aim of this experiment is to compare different optimizers and we need not replicate the "exact" booking curve, we assume the parameters for beta distribution which are calibrated to simulate interspersed arrivals with lower class passengers arriving early in the booking process. Table 3.3 shows the parameters for beta as well as gamma distribution required to fully define the intensity rate.

	Beta Params		Gamma Params					
	α	β	γ	δ	Mean	Variance	Std. Dev	z-Factor
Class 1	4.0	2.0	5.0	3.0	15.0	60.0	7.7	4.0
Class 2	3.0	2.0	6.7	3.0	20.0	80.0	8.9	4.0
Class 3	2.0	3.0	8.3	3.0	25.0	100.0	10.0	4.0
Class 4	2.0	4.0	13.3	3.0	40.0	160.0	12.6	4.0

Table 3.3: Parameters for Gamma Distribution for *z-factor* = 4

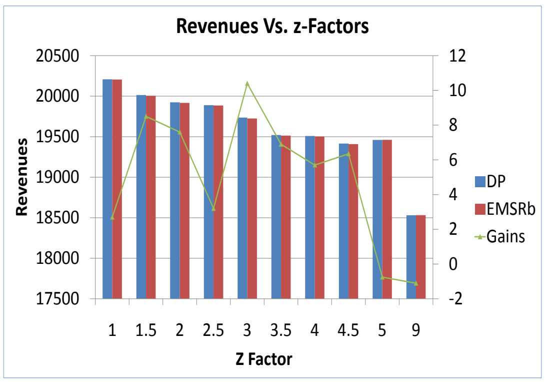

(a) Revenues

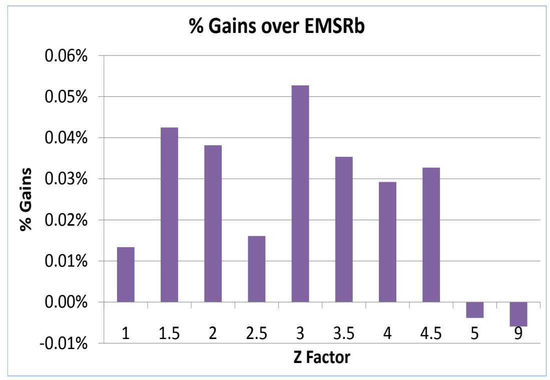

(b) % Gains

Figure 3-8: (a) Revenues for DP and EMSRb (b) % Gains of using DP over EMSRb

Figure 3-8(a) shows the revenues obtained from using DP and EMSRb. We see that the revenues decrease as we increase the demand variance and this phenomenon is independent of the optimizer. This seems intuitive since the introduction of more and more uncertainty will have an adverse effect on any optimizer's performance. We also see from Figure 3-8(b) that the percentage gain of using DP over EMSRb is very small but consistently positive for *z-factor* < 5.0. Table 3.5 shows the load factors obtained by using the two optimizers. An interesting thing to note here is that the load factors obtained by DP are lower than EMSRb up to a *z-factor* of 2.5 and if we further increase demand variance, DP generates higher load factors.

The required percentage difference in the mean revenues obtained by EMSRb and DP to prove one is significantly (statistically) different from the other are shown as a function of confidence level in Table 3.4. We will refer to this table in the following sections to discuss the simulation results. In the above base case, we see that the

Confidence Level	% Difference in Revenues
95%	0.05%
90%	0.04%
85%	0.03%
80%	0.02%
70%	0.01%

Table 3.4: Confidence Levels for various % Changes for N = 1000

result obtained for *z-factor* of 3 is significant at 95% confidence. We also observe that most of the gains obtained by DP(except for *z-factors* of 1 and 2.5) are statistically significant at 85%.

	1	**1.5**	**2**	**2.5**	**3**	**3.5**	**4**	**4.5**	**5**	**9**
DP	94.3	93.2	92.9	91.9	91.7	90.6	90.4	89.9	89.7	86.1
EMSRb	94.9	93.5	93.1	92.2	91.6	90.5	90.2	89.6	89.4	85.4

Table 3.5: Average Load Factors (%) for various z Factors

DP performs slightly better than EMSRb for lower *z-factors* despite slightly lower average load factors. The fare class mix helps us in understanding this observation. Figure 3-9 shows the fare class mix for *z-factor* of 1, 2, 5 and 9 respectively. We can see that for lower *z-factors*, DP accepts more higher class passengers and less lower class passengers than EMSRb whereas for higher *z-factors* DP accepts more lower class passengers and less higher class passengers as compared to EMSRb. This explains that the loss of revenues due to fewer low class passengers is more than offset by the gain in revenues by accpeting more higher class passengers.

We look at the protection levels for Class 1 in Figure 3-10. DP protects more seats for Class 1 in the initial time frames when compared to EMSRb. This trend is independent of the demand variability. We observe that DP is aggressive in protecting seats for higher class passengers. However, we observe that as the day of departure approaches, DP protects less seats than EMSRb for Class 1 passengers. We look at Class 4 availabilities to understand this observation.

Figure 3-11 shows the availability of lowest class 4 seats across time frames for various *z-factors*. DP has less Class 4 availability than EMSRb in initial time frames

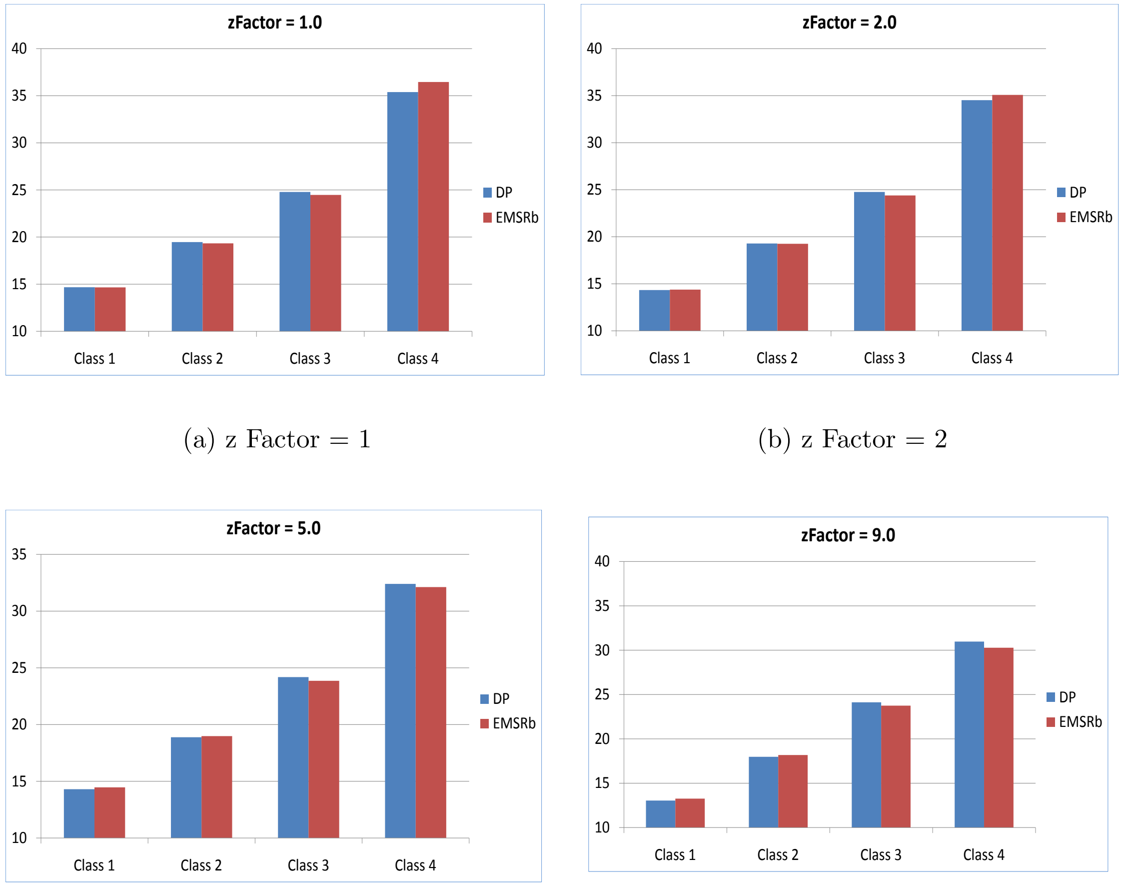

(a) z Factor = 1

(b) z Factor = 2

(c) z Factor = 5

(d) z Factor = 9

Figure 3-9: Fare Class Mix for (a) *z-factor* = 1.0 and (b) *z-factor* = 2.0, (c) *z-factor* = 5.0 and (d) *z-factor* = 9.0

and more availability than EMSRb in later time frames independent of the variability in demand. This observation has a one to one correspondence with our earlier observation about class 1 protection. In the initial time frames, DP is aggressive and protects more seats for Class 1 and hence, less seats are available for Class 4 passengers. In the later time frames, DP protects less seats for Class 1 passengers and this increases the availability for the lowest Class 4 passengers. This observation suggests that DP, because of its Poisson assumption, is more certain about the demand in higher classes (*z-factor* = 1 for Poisson assumption) than EMSRb and hence protects more higher class seats in the initial time frames. As we move later into the booking process, the realized demand has a different variance than assumed

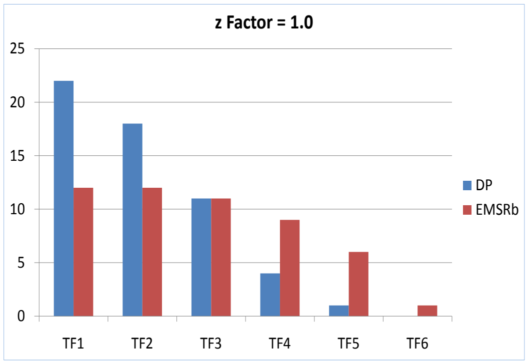

(a) z Factor = 1

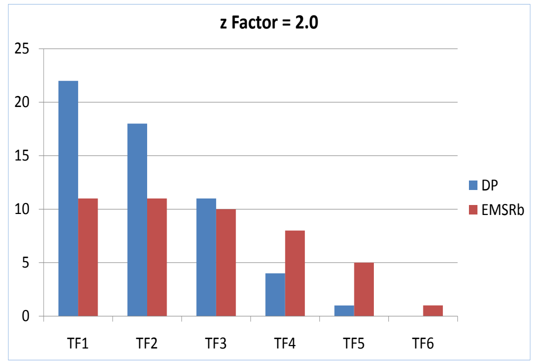

(b) z Factor = 2

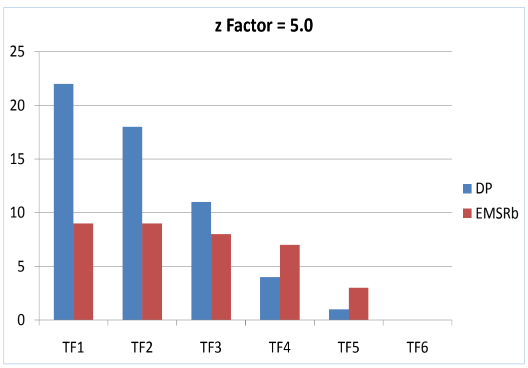

(c) z Factor = 5

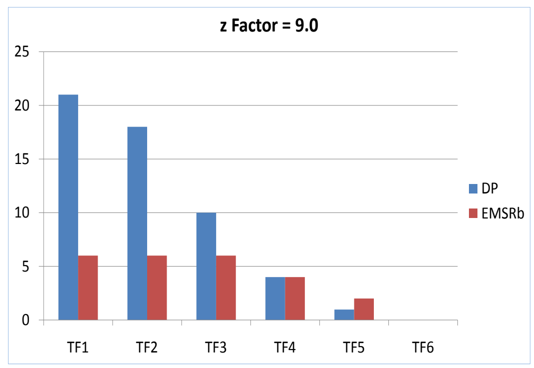

(d) z Factor = 9

Figure 3-10: Class 1 Protection for (a) *z-factor* = 1.0 and (b) *z-factor* = 2.0, (c) *z-factor* = 5.0 and (d) *z-factor* = 9.0

by DP (since simulated *z-factor* is not always 1) through the Poisson assumption. As discussed before, due to passenger characteristics coupled with the restrictions in fare structures, it is empirically observed that the lower class demand (mostly leisure travelers) arrives early in the booking process and the higher class demand (primarily business travelers) arrive later in the booking horizon. The lower availability for the lowest Class 4 seats in DP leads to many Class 4 passengers being rejected in the early part of the booking process. The fewer number of bookings than expected leads DP to increase the availability for lower class passengers which, in turn, means reducing the protection levels for higher class passengers as we have seen in Figures 3-10 and 3-11.

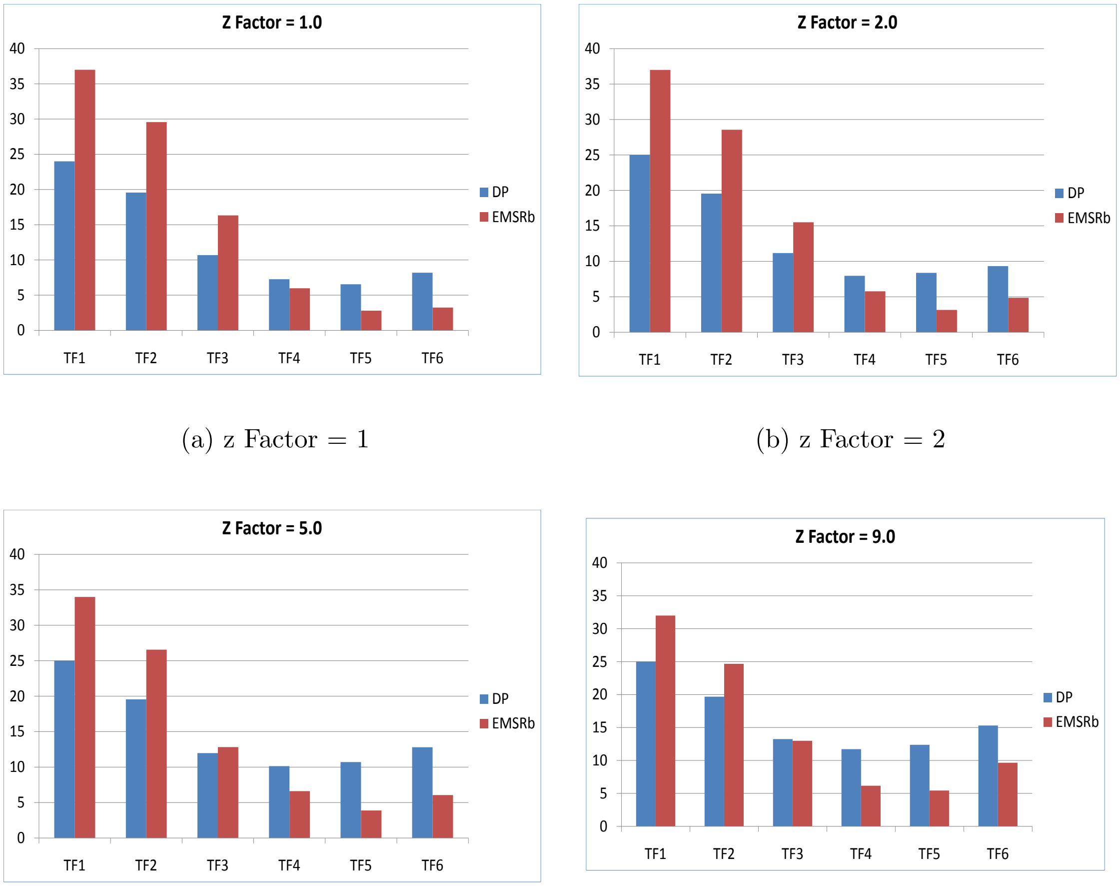

(a) z Factor = 1

(b) z Factor = 2

(c) z Factor = 5

(d) z Factor = 9

Figure 3-11: Class 4 Availability for (a) *z-factor* = 1.0 and (b) *z-factor* = 2.0, (c) *z-factor* = 5.0 and (d) *z-factor* = 9.0

We observe the same phenomenon by looking at Class 4 bookings in Figure 3-12. We observe that as we move across time frames DP falls behind EMSRb in terms of class 4 passengers but around time frame 3 or 4, it increases availability for class 4 and books more passengers in an effort to increase revenues.

Sensitivity Analysis

The results discussed above hold for the inputs as described in Section 3.6. In this section, we discuss the effects of changing inputs on the performance of DP in the single-leg scenario. In particular, we will see the effects of changes in capacity relative to demand, changes in fare ratios and simultaneous change in both capacity and fare

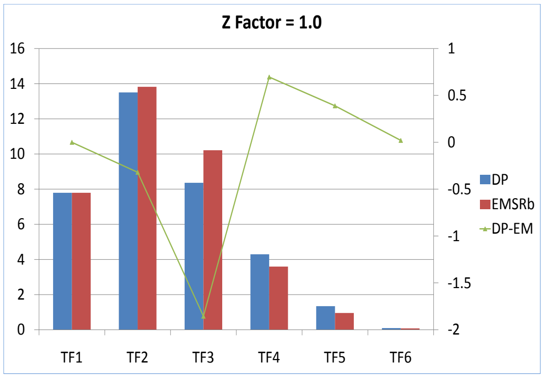

(a) z Factor = 1

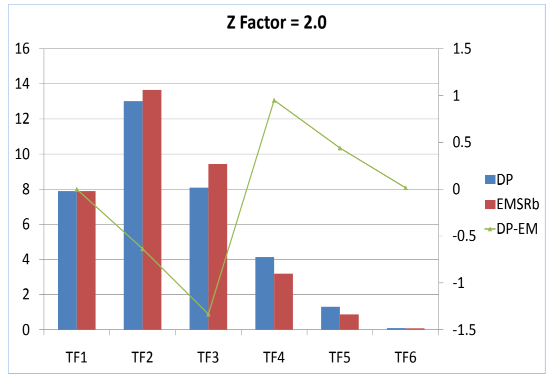

(b) z Factor = 2

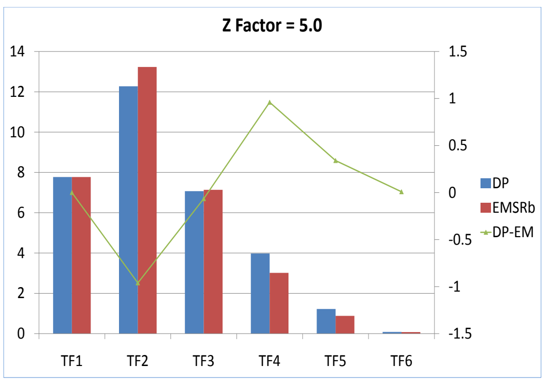

(c) z Factor = 5

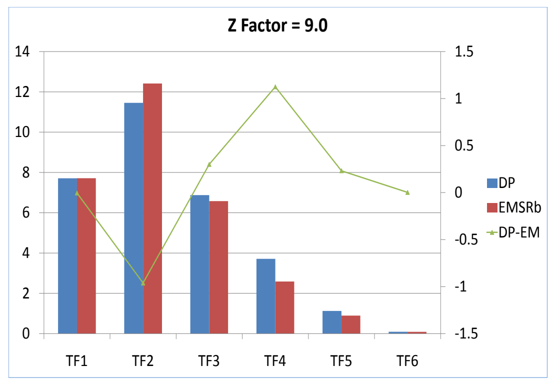

(d) z Factor = 9

Figure 3-12: Class 4 Bookings for (a) *z-factor* = 1.0 and (b) *z-factor* = 2.0, (c) *z-factor* = 5.0 and (d) *z-factor* = 9.0

ratios. The sensitivity to capacity tells us the effects of varying demand levels (since keeping demand fixed while varying capacity is same as varying demand levels while keeping the capacity fixed) and the sensitivity to fare ratios tells us how DP would perform under different fare levels. A sensitivity analysis of this type enables to encompass a wide range of airline markets with varying levels of demand relative to capacity and with varied fare structures.

Capacity Changed to 115

The average load factors obtained from the base case are in the range of 85%–95%. These are very high load factor flight legs that are not typical. Hence, we test the effects of decreasing capacity relative to demand. The results in Figure 3-13 show the impacts of a moderate increase of capacity relative to demand. Recall the significance levels shown in Table 3.4. We observe that the difference in revenues generated by the two optimizers is not statistically significant at 90% for lower *z-factors* but EMSRb outperforms DP for higher *z-factors* and this difference is statistically significant at 95% confidence. We observe the same effect of increasing demand variance as previously discussed – EMSRb outperforms DP as *z-factor* increases.

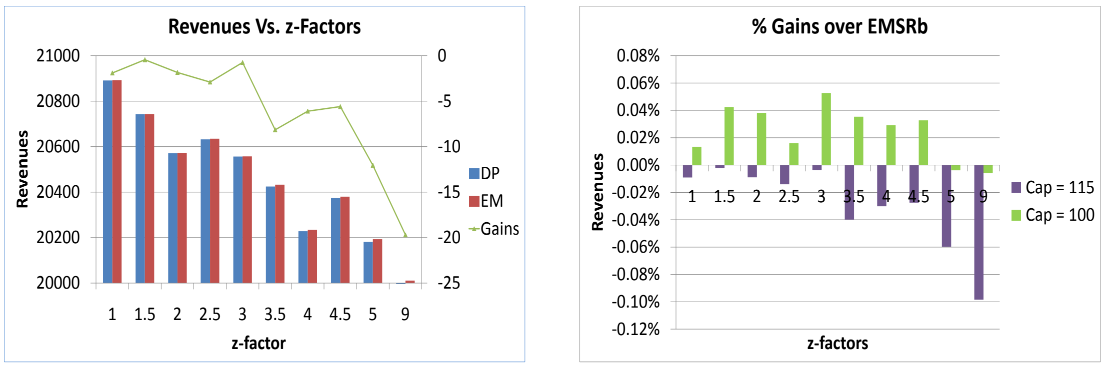

(a) Revenues

(b) % Gains

Figure 3-13: (a) Revenues for DP and EMSRb (Cap = 115) (b) % Gains of using DP over EMSRb (Cap = 115)

We increase the capacity from 100 to 115 and observe that the load factors drop substantially (80%–86%) because of extra capacity. These are much more realistic average leg load factors on a typical high demand flight leg.

	1	**1.5**	**2**	**2.5**	**3**	**3.5**	**4**	**4.5**	**5**	**9**
DP	86.5	85.6	85.2	85.3	85.0	84.7	83.7	84.0	83.1	81.8
EMSRb	86.5	85.7	85.2	85.2	84.9	84.6	83.5	83.8	82.9	81.5

Table 3.6: Average Load Factors (%) for various z Factors (Cap = 115)

Capacity Changed to 130

We further increase the capacity to 130 seats relative to a mean demand of 100. Figure 3-14 shows that EMSRb outperforms DP for all z-factors by a small margin. As discussed before, increasing demand variance adversely affects the performance of DP relative to EMSRb and the differences in revenues are significant at 90% confidence level for *z-factors* greater than 4. We also notice that the average load factors drop to the mid-70s due to an increase in capacity.

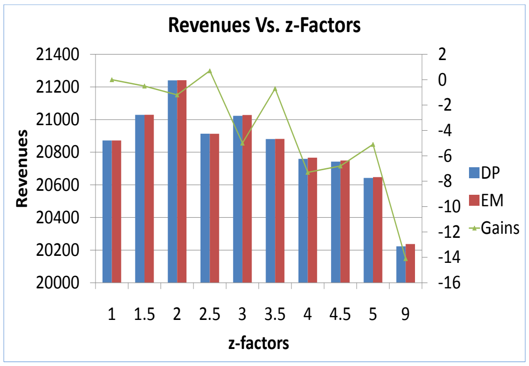

(a) Revenues

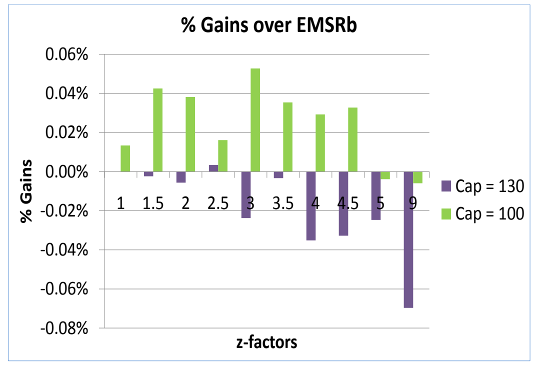

(b) % Gains

Figure 3-14: (a) Revenues for DP and EMSRb (Cap = 130) (b) % Gains of using DP over EMSRb (Cap = 130)

	1	**1.5**	**2**	**2.5**	**3**	**3.5**	**4**	**4.5**	**5**	**9**
DP	76.3	77.1	77.6	76.4	77.0	76.4	76.1	76.3	76.0	74.4
EMSRb	76.3	77.1	77.6	76.4	76.9	76.4	76.0	76.2	75.9	74.2

Table 3.7: Load Factors (%) for various z Factors (Cap = 130)

Fare Ratio Decreases

In the previous two sensitivity tests, we change only the capacity relative to demand. In this experiment, we keep the capacity equal to 100 seats for a mean demand of 100 but we change the fare ratios as shown in Table 3.8.

	Original Fares ($)	**New Fares ($)**
Class 1	400	250
Class 2	300	200
Class 3	200	150
Class 4	100	100

Table 3.8: Modified Fare Levels: Low Fare Ratios

We note that the fare ratios in the modified fare levels are lower than the original fare ratios. The highest to the lowest fare ratio in the original fare structure 4:1 where as it is 2.5:1 under the modified fare levels. The new fare levels represent a market in which the fares are very compressed. Figure 3-15 shows that DP performs better than EMSRb and percentage gain of using DP over EMSRb increases with increasing demand variance. The results are statistically significant at a confidence level of 95% for *z-factors* greater than 2.

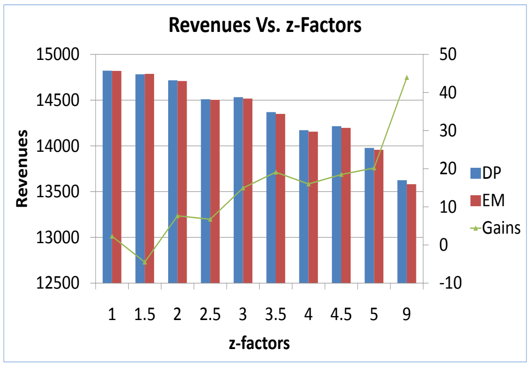

(a) Revenues

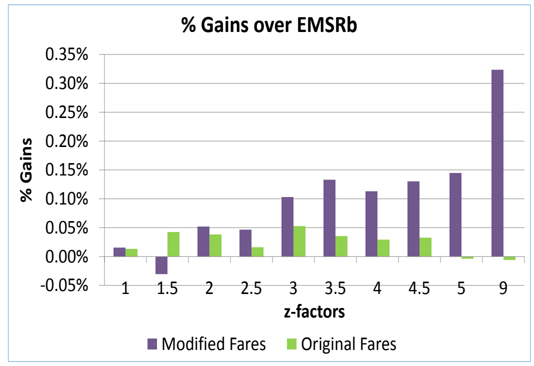

(b) % Gains

Figure 3-15: (a) Revenues for DP and EMSRb (Low Fare Ratios) (b) % Gains of using DP over EMSRb (Low Fare Ratios)

Under the new fare structure, the average load factors for DP are always lower than EMSRb (3.9).

	1	**1.5**	**2**	**2.5**	**3**	**3.5**	**4**	**4.5**	**5**	**9**
DP	94.6	94.1	93.8	92.4	92.3	91.3	90.5	90.4	89.2	86.5
EMSRb	95.2	94.7	94.4	93.1	93.1	91.7	91.1	90.9	89.7	87.3

Table 3.9: Average Load Factors (%) for various z Factors (Low Fare Ratios)

Fare Ratio Increases

Table 3.10 compares the new fare levels to the old ones. We note that the fare ratios in the modified fare levels are higher than the original fare ratios. The highest to the lowest fare ratio in the original fare structure 4:1 where as it is 7:1 under the modified fare levels. The new fare levels represent a market in which the fare levels are quite dispersed. We observe from Figure 3-16 that EMSRb outperforms DP independent of the demand variance. Table 3.11 shows that the load factors for DP are always

	Original Fares ($)	New Fares ($)
Class 1	400	700
Class 2	300	500
Class 3	200	300
Class 4	100	100

Table 3.10: Modified Fare Levels: High Fare Ratios

lower than EMSRb which causes a loss in revenues. We also note that the results are statistically significant at 95% confidence levels for *z-factors* greater than 2.

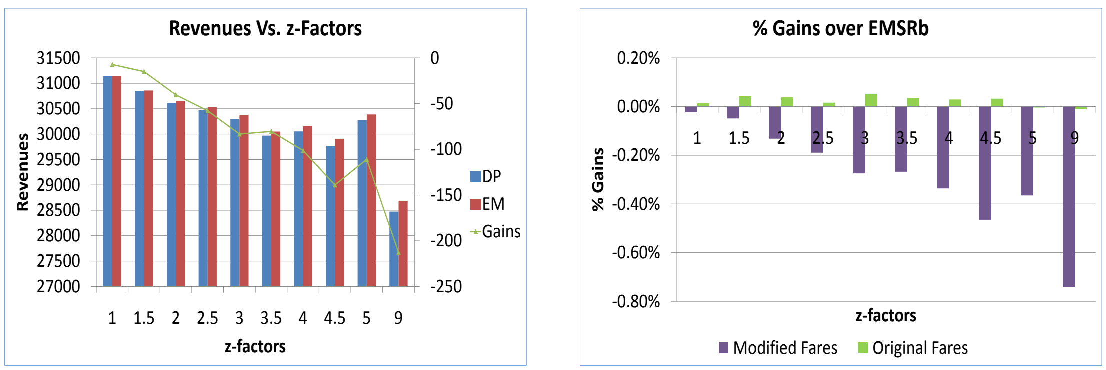

(a) Revenues

(b) % Gains

Figure 3-16: (a) Revenues for DP and EMSRb (High Fare Ratios) (b) % Gains of using DP over EMSRb (High Fare Ratios)

	1	**1.5**	**2**	**2.5**	**3**	**3.5**	**4**	**4.5**	**5**	**9**
DP	93.8	93.3	92.2	91.8	91.1	90.2	89.9	89.3	89.4	85.3
EMSRb	93.6	93.2	91.7	90.9	89.8	88.6	88.0	87.5	87.3	81.5

Table 3.11: Average Load Factors (%) for various z Factors (High Fare Ratios)

We observe that DP outperforms EMSRb when the fare ratio decreases and under performs relative to EMSRb when the fare ratio increases. Table 3.12 shows the availability for the lowest class (Class 4) for both DP and EMSRb under various fare ratios. We observe that as fare ratio gets smaller, EMSRb makes more Class 4 seats available and there is no significant change in DP Class 4 availability. This, in turn, means that EMSRb protects less seats for Class 1. We observe the exact opposite effect when the fare ratios become larger – EMSRb makes less Class 4 seats available and this, in turn, means more seats are protected for Class 1. This is intuitive since EMSRb takes into account the fare ratios and as the fares get lower, the highest class seat is not as valuable as it was before and hence EMSRb protects less for the highest class and makes more Class 4 seats available. We observe the exact same phenomenon as discussed above if we look at the fare class mix for a typical airline *z-factor* of 5.0 for the three fare ratios as shown in Figure 3-17.

	Class 1 Protection		Class 4 Availability	
Fare Ratio	DP	EMSRb	DP	EMSRb
Low	22	7	26	42
Base	22	9	25	34
High	23	10	22	26

Table 3.12: Class 1 Protection and Class 4 Booking Limits at TF1 for *z-factor* = 5

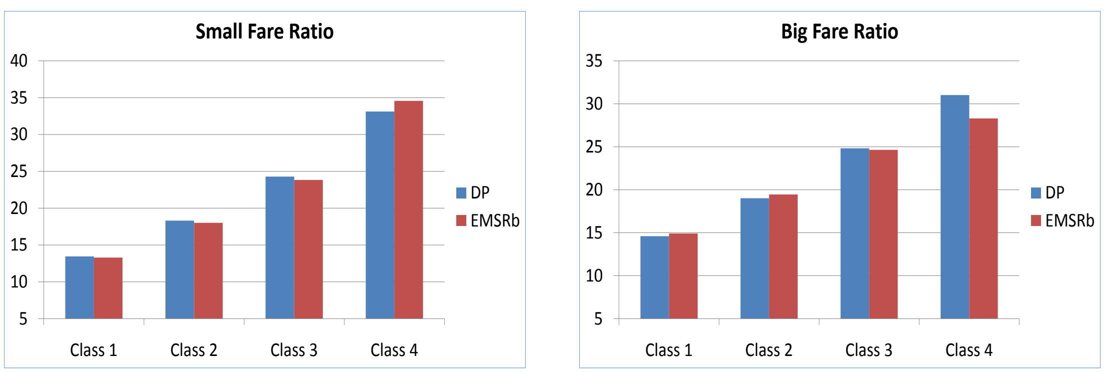

(a) Small Fare Ratio (b) Big Fare Ratio

Figure 3-17: Fare Class Mix for (a) Small Fare Ratio (b) Big Fare Ratio

<u>*Joint Sensitivity Analysis*</u>

The results discussed previously represent a marginal sensitivity analysis in which an individual variable was varied while keeping the others constant. In the joint sensitivity analysis, we vary both the capacity and the fare ratios. Figure 3-18 shows that the effect of fare ratio is dominant and fare ratio is important independent of the capacity being used. DP always performs well under low fare ratios and performs worse than EMSRb under high fare ratios.

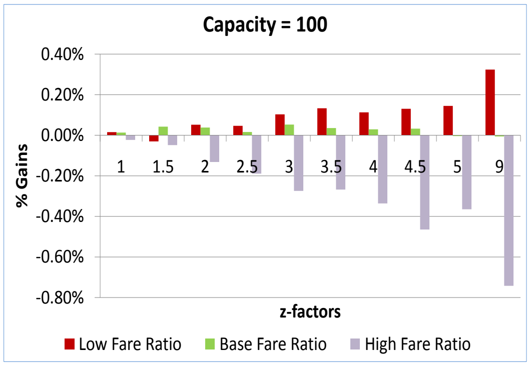

(a) Revenues

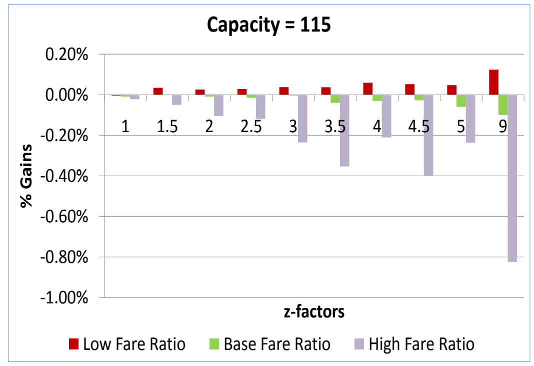

(b) % Gains

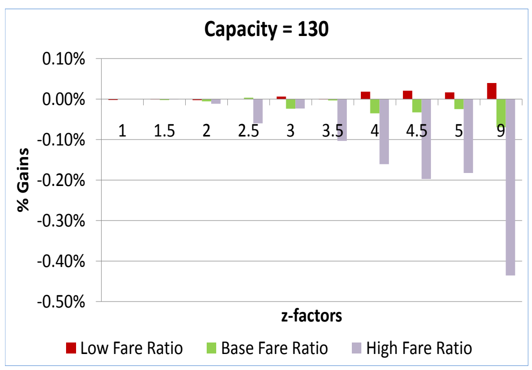

(c) % Gains

Figure 3-18: % Gains over EMSRb (a) Cap = 100 (b) Cap = 115 and (c) Cap = 130

To summarize, we observe that in the base case DP produces higher revenues than EMSRb for *z-factors* < 5, though the gains were very small (0.01%-0.05%) with our base case set-up. We recall from Table 3.4 that gains less than 0.05% are not significant at 95% confidence level. The higher revenues obtained by DP can be explained by two phenomena in different *z-factor* regimes. For $1 \leq$ *z-factor* ≤ 2.5, DP produces lower load factors but accepts more class 1, 2, 3 bookings and less class 4 bookings. This fare class mix gives a boost to the yield generated through this policy and hence, despite lower load factors, DP modestly outperforms EMSRb. For *z-factor* > 2.5, DP accepts more lower class bookings and less higher class bookings but it generates higher load factors than EMSRb and again, modestly outperforms EMSRb. We also observe that with increasing variance in the demand process, DP accepts more Class 4 passengers and fewer Class 1 passengers. In the initial time frames, DP makes less seats available for Class 4 as compared to EMSRb but in the later time frames it makes more seats available than EMSRb.

In all the experiments we did, we observe that for higher *z-factors*, EMSRb outperforms DP. The effects of higher realized demand variance than assumed in DP outweighs the benefits of a dynamic optimization in DP. The revenue differences between DP and EMSRb, although some of them are small and significant only at 80% confidence level, provide some key insights into the performances of DP and EMSRb as discussed above. The sensitivity analysis shows us that DP does better in cases when the capacity relative to demand is high and when the fare ratios are low. We showed the effects of changing fare ratios on the performance of EMSRb and DP and the fare class mix verified our explanation on the effect of fare ratios on the performance of these optimizers. The experiment on joint sensitivities reveal that fare ratios dominate the effect of capacity change when we simultaneously change both of these parameters and this can partly be seen by the magnitude of the marginal sensitivities of each of the variables.

3.7.2 Interspersed Arrivals - Uniform Distribution

In the previous experiment, we assumed that the probability of a lower fare-class arrival is higher in initial time frames than in later time frames and the probability of a higher fare-class arrival is lower in initial time frames than in later time frames. We specify this through the beta distribution in the intensity function. In this experiment, we study the effect of relaxing the traditional assumption of lower fare class arriving first in the interspersed arrival pattern. For this experiment, we assume that the probability of arrival of any fare class in any time frame is uniformly distributed. To model this, we need to specify the correct parameters for the beta distribution. Next, we discuss the modification necessary to the beta distribution parameters to specify the above-described arrival pattern.

The probability density function (pdf) of the beta distribution is given by

$$f(x) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha,\beta)}$$

B, defined as the beta function, is a normalization constant to ensure that the total probability integrates to one. The beta function (B) is defined in terms of the gamma function (Γ)as follows:

$$B(\alpha,\beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}$$

If we choose $\alpha = \beta = 1$ and use $\Gamma(n) = (n-1)!$, we see that this pdf equals to one which represents the uniform distribution. Hence, beta distribution with shape parameters (1,1) is used to simulate uniform distribution.

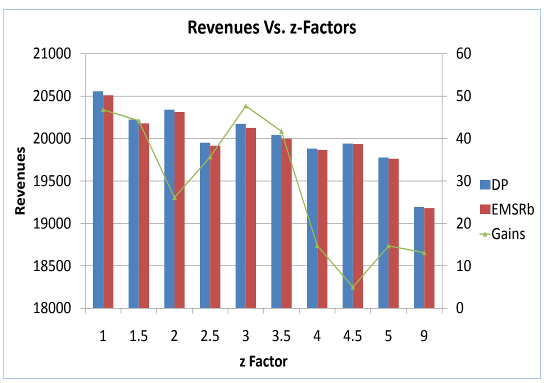

(a) Revenues

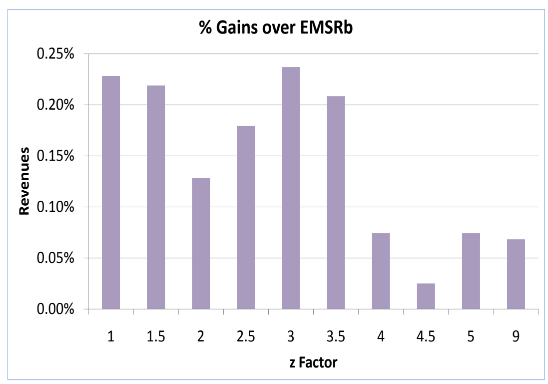

(b) % Gains

Figure 3-19: (a) Revenues for DP and EMSRb (b) % Gains of using DP over EMSRb

Figure 3-19(a) shows the revenues obtained by DP and EMSRb. We observe that, as in the previous case, the revenues obtained by both the optimizers decrease as we increase the demand variance which is intuitively correct. The benefits of using DP over EMSRb increase as we relax the assumption on arrival pattern. We get this result because, even within a time frame, DP is a dynamic optimizer since it produces booking limits (or bidprices) as a function of remaining capacity and bookings in hand, while EMSRb is a static optimizer that assumes a predetermined arrival sequence with lower classes arriving earlier and higher class demand arriving later in the booking process. In this experiment, the actual demand arrival process does not conform to EMSRb's assumptions and hence, EMSRb performs slightly worse than DP. Figure 3-19(b) shows that the % gains of DP over EMSRb is consistent through the range of *z-factors* that we simulated and the benefits of using DP are larger than in previous experiment. We also note that, except for *z-factor* equal to 4, the gains are statistically significant at 95% confidence level. Table 3.13 shows that the load factors obtained by DP are always higher than EMSRb. We will look into the fare class mix to gain deeper insights into this observation.

	1	**1.5**	**2**	**2.5**	**3**	**3.5**	**4**	**4.5**	**5**	**9**
DP	96.1	94.6	94.4	93.1	93.3	92.3	92.1	91.4	91.0	87.9
EMSRb	95.4	94.1	93.8	92.4	92.6	91.6	91.4	90.7	90.3	87.0

Table 3.13: Average Load Factors (%) for various z Factors

We have seen that DP has higher load factor than EMSRb for all the simulated *z factors*. Figure 3-20 shows that DP accepts more lower class passengers and fewer higher class passengers. The revenues gained by accepting more lower class passengers compensates for the revenues lost by rejecting higher class passengers for all *z factors*. As we have seen before, DP is a bit too aggressive in the initial time frames and makes fewer seats available for Class 4 than EMSRb. In the later time frames, however, it makes more seats available in class 4 than EMSRb as shown in figure 3-21.

We observe larger benefits from DP if we assume that passenger arrival process is interspersed arrivals with uniform distribution. We also show that the underlying mechanics of each of the optimizers is the same as before (DP is aggressive in initial time frames – protects more seats than EMSRb for Class 1 and fewer seats are made available for Class 4. However, as we move across the timeframes, DP protects less Class 1 and makes more seats available for Class 4 as compared to EMSRb). The next step is to study the stability of these results when the input parameters are varied.

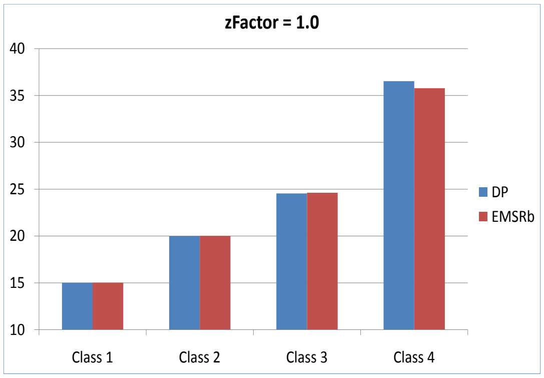

(a) z Factor = 1

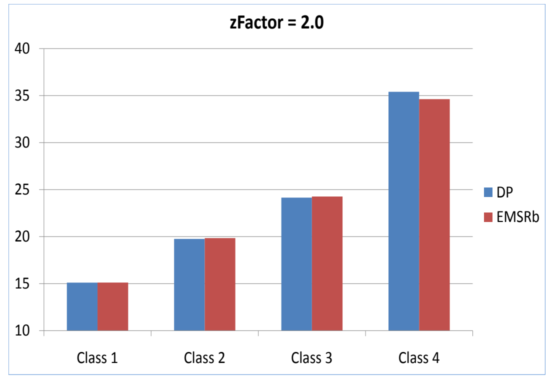

(b) z Factor = 2

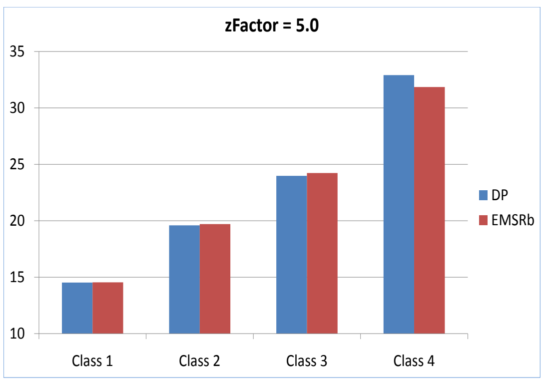

(c) z Factor = 5

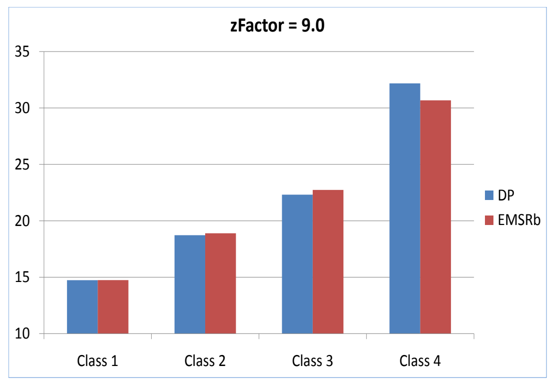

(d) z Factor = 9

Figure 3-20: Fare Class Mix for (a) *z-factor* = 1.0 and (b) *z-factor* = 2.0, (c) *z-factor* = 5.0 and (d) *z-factor* = 9.0

Sensitivity Analysis

We do the sensitivity analysis for the case with uniform arrivals in the same way as we did for the case when the arrivals are interspersed with lower class arriving early. We test these sensitivities for a change in capacity and and a change in fare ratios.

Capacity Increased to 115

We observe from Figure 3-22 that the gains of using DP are much less than those with a capacity of 100. We observed the same effect of an increase in capacity in the previous case. DP moderately outperforms EMSRb for lower *z-factors* and performs worse than EMSRb for higher *z-factors*(> 5). The load factors obtained by DP are

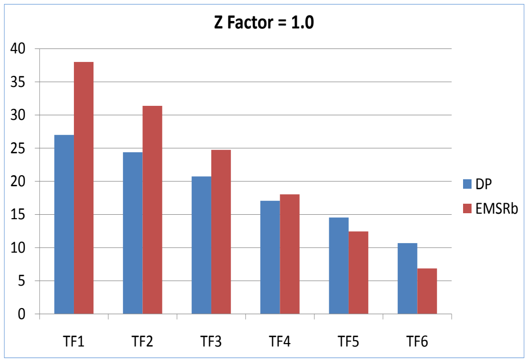

(a) z Factor = 1

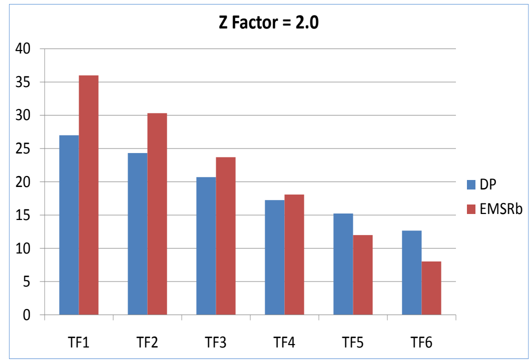

(b) z Factor = 2

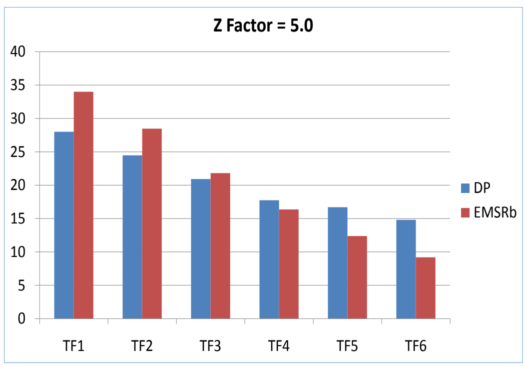

(c) z Factor = 5

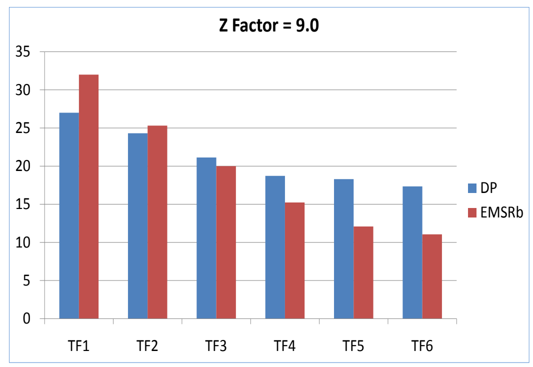

(d) z Factor = 9

Figure 3-21: Class 4 Availability for (a) *z factors* = 1.00 and (b) *z factors* = 2.00, (c) *z factors* = 5.00 and (d) *z factors* = 9.00

always higher than EMSRb as shown in Table 3.14.

	1	**1.5**	**2**	**2.5**	**3**	**3.5**	**4**	**4.5**	**5**	**9**
DP	86.9	86.2	86.3	85.6	85.3	84.2	85.1	84.6	84.1	80.8
EMSRb	86.8	86.1	86.2	85.5	85.1	84.0	84.8	84.4	83.9	80.5

Table 3.14: Average Load Factors (%) for various z Factors (Cap = 115)

Capacity Increased to 130

We observed that as we increase the capacity from 100 to 115, the gains of DP over EMSRb decrease. If we increase the capacity further to 130, almost all the gains of using DP vanish as shown in Figure 3-23. DP generates negligibly small

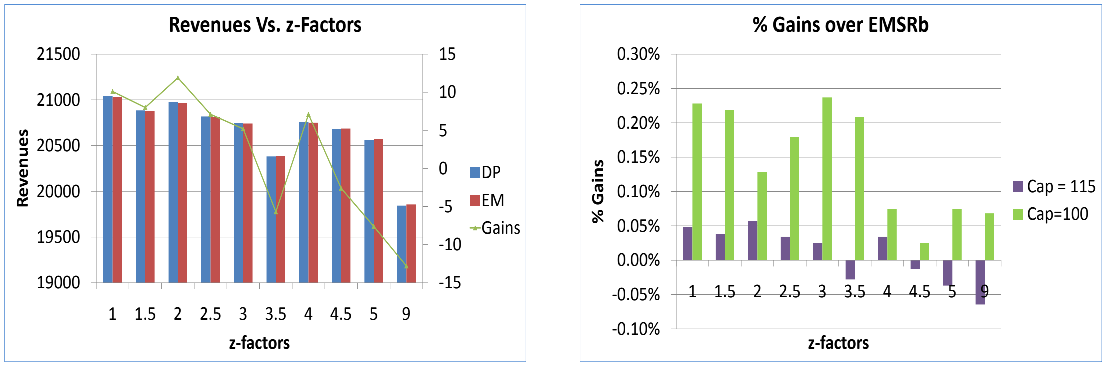

(a) Revenues

(b) % Gains

Figure 3-22: (a) Revenues for DP and EMSRb (Cap = 115) (b) % Gains of using DP over EMSRb (Cap = 115)

incremental revenues over EMSRb for small *z-factors* and DP under performs for higher *z-factors*. These results of effects of capacity increase are consistent with our results from previous case.

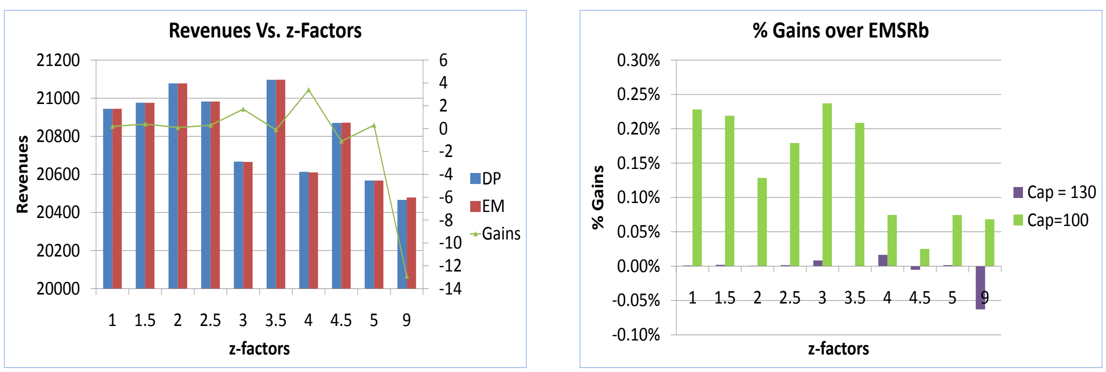

(a) Revenues

(b) % Gains

Figure 3-23: (a) Revenues for DP and EMSRb (Cap = 130) (b) % Gains of using DP over EMSRb (Cap = 130)

Fare Ratio Decreased

We modify the fare ratios in similar way as we show in Table 3.16.

In this case, as seen in the previous experiment, DP performs better than EMSRb and percentage gain of using DP increases with increasing demand variance as shown

	1	1.5	2	2.5	3	3.5	4	4.5	5	9
DP	73.9	72.8	72.6	71.6	71.8	70.9	70.8	70.3	70.0	67.6
EMSRb	73.4	72.3	72.2	71.0	71.2	70.5	70.3	69.8	69.5	66.9

Table 3.15: Average Load Factors (%) for various z Factors (Cap = 130)

	Original Fares ($)	New Fares ($)
Class 1	400	250
Class 2	300	200
Class 3	200	150
Class 4	100	100

Table 3.16: Modified Fare Levels: Low Fare Ratios

in Figure 3-24. We also observe that DP performs worse than original fares for smaller *z-factors* and performs better than the original fares for higher *z-factors*. Table 3.17 reports the percentage change in total bookings for each class due to a shift from the base fare structure to the current fare structure for two *z-factors*. We observe that the for small *z-factors*, a shift to the compressed fare structure decreases bookings in the higher classes and increases bookings in the lower classes where as for higher *z-factors*, we observe that the number of bookings decrease only in fare class 2 and increase in all other fare classes. The important thing to observe here is that, unlike for small *z-factors* where the change is almost parallel between DP and EMSRb, for higher *z-factors*, the benefits of moving to a compressed fare structure are more in DP than EMSRb for higher classes. The extra revenues generated by booking more higher class passengers more than offsets the lost revenues due to rejecting lowest class 4 passengers. As we have seen before, EMSRb increases the availability for lowest Class 4 as the fare ratios get lower. We observe the same phenomenon here for both the cases of demand variability. Since for small demand variability, both EMSRb and DP lose higher class passengers although EMSRb attracts more Class 4 passengers, the difference between DP and EMSRb revenues decrease from the base case. This same effect is also present for higher *z-factors* but it is accompanied by DP capturing more higher class passengers which offsets the lost revenues due to smaller availability in lowest Class 4.

Table 3.18 reports the load factors obtained by DP and EMSRb. We observe that the load factors of DP are always higher than that of EMSRb as we have seen in the previous experiment.

	z-factor=1.0		*z-factor*=9.0	
Class	DP	EMSRb	DP	EMSRb
1	-1.41%	-1.45%	5.18%	3.81%
2	-1.10%	-1.24%	-0.46%	-2.26%
3	-1.21%	-1.47%	2.23%	0.24%
4	1.32%	2.79%	0.37%	5.85%

Table 3.17: Percentage Change in Number of Bookings due to Shift From Base Fare Ratio

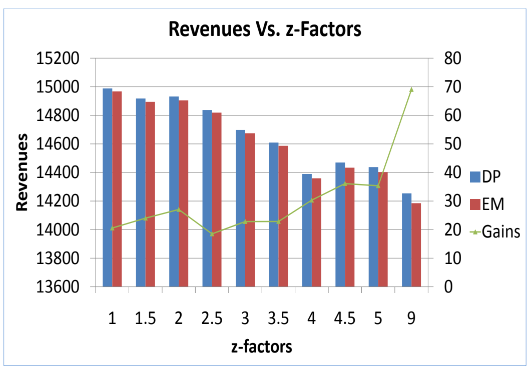

(a) Revenues

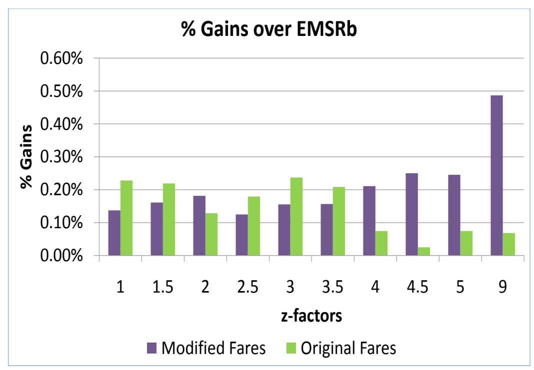

(b) % Gains

Figure 3-24: (a) Revenues for DP and EMSRb (Low Fare Ratios) (b) % Gains of using DP over EMSRb (Low Fare Ratios)

	1	**1.5**	**2**	**2.5**	**3**	**3.5**	**4**	**4.5**	**5**	**9**
DP	95.8	95.1	94.9	94.2	93.2	92.6	91.4	91.4	90.9	89.2
EMSRb	95.6	94.8	94.6	94.0	92.9	92.4	91.2	91.2	90.7	89.0

Table 3.18: Average Load Factors (%) for various z Factors (Low Fare Ratios)

Fare Ratio Increased

We modify the fare ratios in similar way as we show in Table 3.19.

	Original Fares ($)	New Fares ($)
Class 1	400	700
Class 2	300	500
Class 3	200	300
Class 4	100	100

Table 3.19: Modified Fare Levels: High Fare Ratios

We observe that for lower *z-factors*, DP outperforms EMSRb and the trend reverses for *z-factors* greater than 2.5. DP outperforms EMSRb for all *z-factors* in the original fare structure. These results are consistent with our previous sensitivity analyses on fare structures. Table 3.20 shows the load factors which are qualitatively the same as we had in the base case – DP has higher load factors than EMSRb for all *z-factors*.

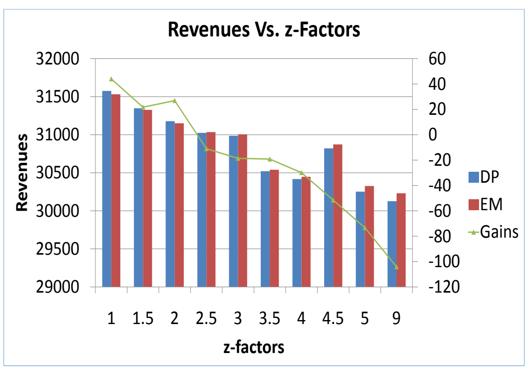

(a) Revenues

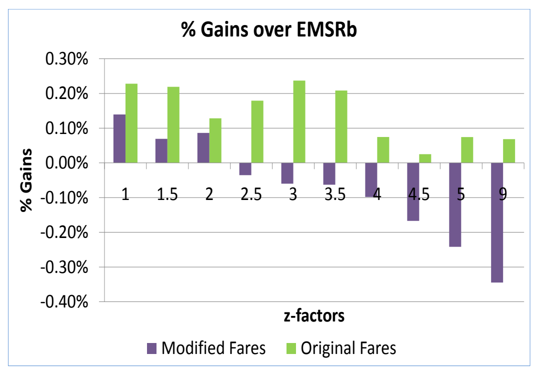

(b) % Gains

Figure 3-25: (a) Revenues for DP and EMSRb (High Fare Ratios) (b) % Gains of using DP over EMSRb (High Fare Ratios)

	1	1.5	2	2.5	3	3.5	4	4.5	5	9
DP	96.0	94.8	94.3	93.4	93.3	91.9	91.2	91.5	90.9	87.8
EMSRb	95.1	93.9	93.1	92.4	91.9	90.6	89.9	89.9	89.5	85.6

Table 3.20: Average Load Factors (%) for various z Factors (High Fare Ratios)

Joint Sensitivity Analysis

The joint sensitivity analysis again shows us that DP performs well under low fare ratios and performs worse than EMSRb under high fare ratios. These results are consistent with our previous joint sensitivity analysis. This analysis also shows us the importance of demand factor[2] — In our experiments, DP performs very well when demand factor is unity but lowering demand factor (increasing capacity) affects the performance of DP adversely. The results are shown in Figure 3-26.

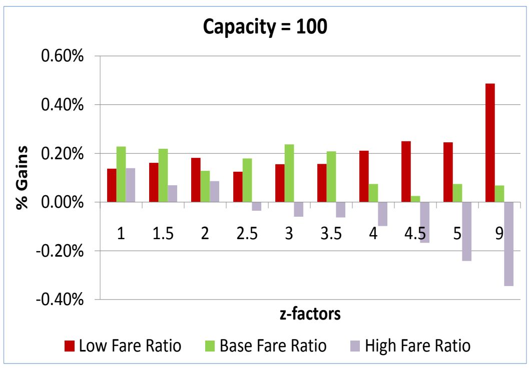

(a) Revenues

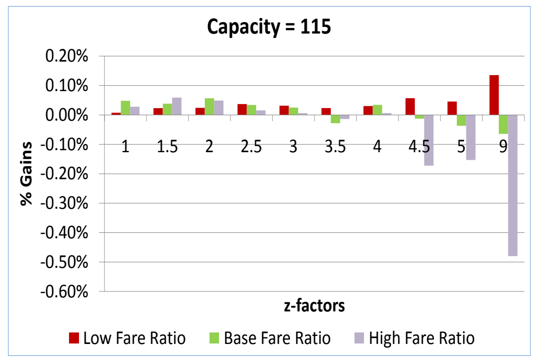

(b) % Gains

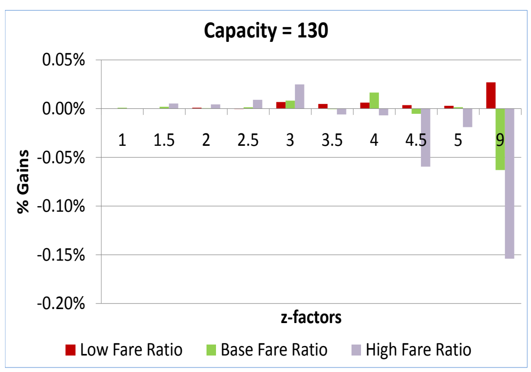

(c) % Gains

Figure 3-26: % Gains over EMSRb (a)(Cap = 100) (b) (Cap = 115) and (c) (Cap = 130)

To summarize, DP consistently produces higher revenues than EMSRb and the observed gains were larger (0.05%-0.20%) than before when we assume the traditional arrival pattern. These gains are statistically significant at 95% confidence level. The higher revenue obtained by DP can be attributed to high load factors it generated

[2]Demand Factor is defined as the ratio of demand to capacity

independent of the *z-factor*. The gains in revenue due to higher class 4 bookings more than compensated the loss of revenues due to rejection of higher-fare class passengers. With increasing variance in the demand process, DP accepts more Class 4 passengers and fewer Class 1 passengers as compared to EMSRb as seen before. DPs protection policy for class 4 is same as before — DP protects less seats for Class 4 in initial TFs but in the later TFs it protects more when compared with EMSRb. The sensitivity analysis shows us the same qualitative results as we observed in the previous case that low fare ratios and low demand factors favor DP along with smaller demand variability.

3.8 Summary

We built a simulator to observe the effects of various levels of demand variability on the performance of DP. We simulated a single flight leg with two basic optimizers — EMSRb and DP. This simulation environment excluded the effects of competition as well as that of forecasting accuracy and it matches the assumptions of DP models for the simulated *z-factor* of 1. A sensitivity analysis on various assumptions of demand factors as well as fare levels was done for two different types of passenger arrival pattern. We summarize our main findings from this experiment in this section.

The difference in revenues generated by DP and EMSRb is small for our test data. We found that this difference is usually in the 0.01%–0.5% in all our experiments. We studied the effects of following factors in a single flight leg example:

1. Actual Demand Variance
2. Fare Ratios
3. Capacity (or Demand Factor)
4. Arrival Pattern (Front-Loaded or Uniform Interspersed Arrivals)

We found that DP performs better than EMSRb for the baseline case (Demand Factor = 1.0 and base fare ratio) except when $z - factors > 5$, which is actually more typical for airline RM systems. Our results show that, when DP outperforms EMSRb, it can be due to either higher load factor or better fare-class mix (higher yield) despite lower load factor. A common theme across all the simulation results was that DP always has lower availability for lower classes (and more protection for higher classes) in the initial time frames and has more availability for lower classes (less protection for higher classes) in later time frames when compared to EMSRb. We also conclude that the effect of changing fare ratios is dominant when compared to changes in demand factor. Based on our experiments, we summarize that DP outperforms EMSRb if the demand variability is low, if load factors are high ($> 90\%$), if fare ratios are low, or if the arrival pattern is uniformly distributed across all time frames for all fare classes.

We discussed in the preceding paragraphs that actual demand variance is a critical factor in determining the performance of DP relative to EMSRb. MIT's airline data project website reports domestic US carriers' average load factor was 81% in 2008. Most airline leg load factors are in the range of 75%–85% but can range as low as 60% and as high as 95%, as gathered after discussion with the members in the PODS consortium. Most of the airlines have a forecasting system that generates forecasts of remaining demand prior to departure by fare class. Figure 3-27 shows *z-factors*[3] of fare class forecasts as a function of remaining days to departure as reported by one of the PODS Consortium airlines[4]. In this study, the *z-factors* were calculated by considering total remaining demand at each timeframe. We observe that typical *z-factors* of fare class forecasts are between 4 and 5. We should keep these in mind while interpreting the results obtained in this chapter.

We revisit our results for *z-factors* ranging between 4 and 5. We find that for more realistic *z-factors*, there is no significant difference between the performance of DP relative to EMSRb in the base case of this simulation study. The increase in capacity relative to demand (resulting in lower load factors) or an increase in fare ratios adversely affects DP's performance while a decrease in fare ratios leads DP to outperform EMSRb.

[3] *z-factor* $= \frac{\sigma^2}{\mu}$

[4] The plot was used with due consent from the airline

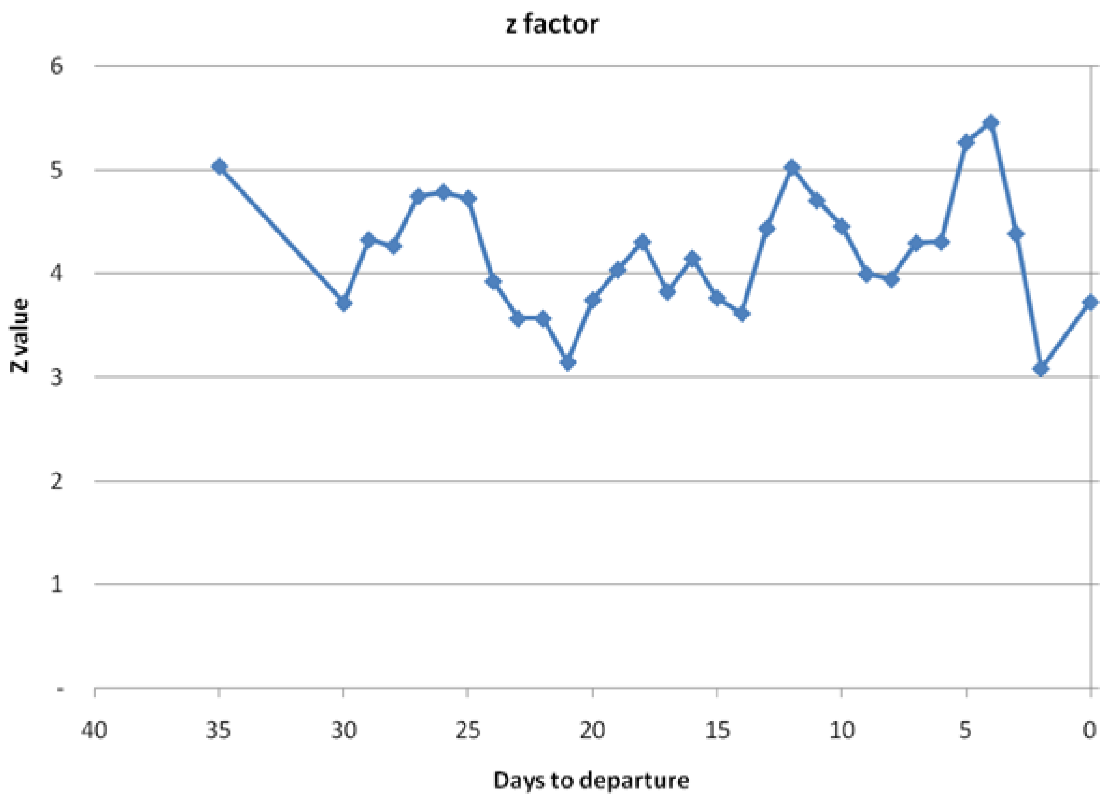

Figure 3-27: Typical Airline *z-factors*

We conclude this discussion by noting that high demand variance affects the performance of DP adversely and the benefits of having a dynamic booking limit policy (as in DP) are negated by the fact that the assumed demand variance in DP models is not equal to the realized variance.

Chapter 4

Passenger Origin Destination Simulator

The simulator developed in the previous chapter represents a test-bed to evaluate the performance of an optimizer under ideal conditions in which the simulator matches the assumptions of the optimization models. In this chapter, we describe the Passenger Origin Destination Simulator (PODS), which is a more realistic simulator. PODS incorporates the elements of forecasting and competition between airlines in a network, both of which were absent from our simulator described in Chapter 3. PODS will also enable us to compare the optimization models in much bigger networks as we will see in the next chapter. We also introduce two network RM methods using DP which are an extension of the standard leg-based DP models.

4.1 PODS Background

PODS evolved from the Decision Window Model (DWM) developed by Boeing to model the passenger choice of selecting different flight alternatives (based on departure time preference, path preference preference, airline preference etc.) for reaching their destination from their origin[1]. The Passenger Origin Destination Simulator

[1]Boeing Commercial Airplane Group, *Decision Window Path Preference Methodology Time Mode*, November, 1993

(PODS) was developed by Hopperstad, Berge and Filipowski in 1997 at the Boeing Company (Hopperstad, 2005) to model the airline booking process with competing airlines trying to maximize passenger revenues in different competitive network configurations. PODS has been further developed by the PODS Consortium, a research alliance between MIT and 9 leading international airlines.

4.2 PODS Architecture

The central idea of PODS is to simulate the interactions between the passengers and the airline RM systems. At one end of the spectrum, the passengers seek bookings for their travel and they choose from a multitude of fares, paths[2], and airlines serving their places of interest. On the other end, the airlines decide the fare structure for each Origin-Destination (OD) market and control the number of seats available for each fare class on every flight serving each of these markets. The PODS RM architecture is similar to that of a typical third generation RM system as was shown in Figure 2-1. Figure 4-1 shows the airline booking process is modeled as an interaction between two components in PODS – the Passenger Choice Model and the RM system. Before we discuss these two components, we present the basic terminology of simulation in PODS.

In PODS, each departure day is called a "sample" and one "trial" consists of simulating 600 "samples" on each leg and one "run" of the simulator consists of 5 independent "trials". In order to start a PODS run, user-defined inputs must be fed into the system. As the simulation progresses, these values are eventually replaced by simulated values. We discard the first 200 samples in each trial to eliminate the effects of the starting parameter values. The overall result for each simulated airline is averaged across the last 400 samples of 5 trials giving 2000 daily simulations to ensure statistical significance of simulation results.

For each sample, the booking process in PODS begins 63 days before departure

[2] A path is defined as a collection of flight legs that form an itinerary between an origin and a destination within a network

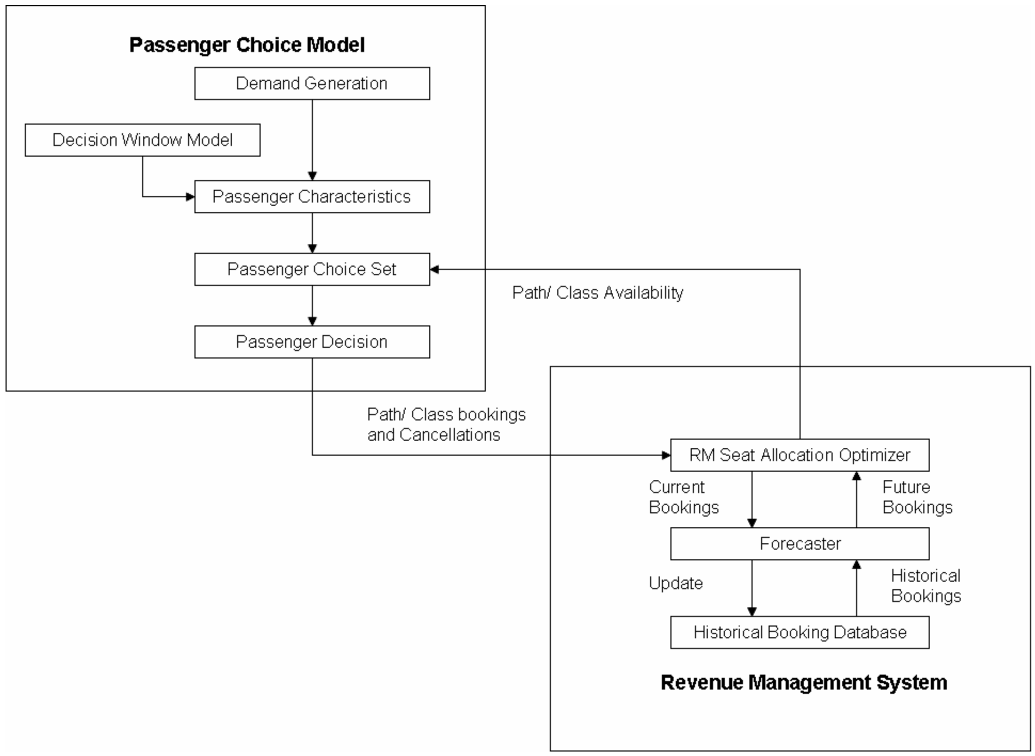

Figure 4-1: PODS Architecture (Source: The PODS Primer)

and is divided into 16 time frames that end on the departure day. At the beginning of the booking process, the time frames are wider due to less anticipated booking activity but as the departure day draws nearer, the time frames become smaller in anticipation of increased booking activity. The RM systems of the airlines update the path/class[3] availabilities at the start of each time frame, while passenger events such as bookings and cancellations occur randomly within each time frame. In this thesis, we will use the terms path/class or ODF interchangeably.

4.2.1 Passenger Choice Model

The capability to model passengers' behavior is an important characteristic that differentiates PODS from other RM simulators. The Passenger Choice Model in PODS models the behavior of passengers and their choices relative to path/class availability

[3]A path is defined as a sequence of flight legs that form an itinerary between an origin and a destination within a network. There are multiple fare classes offered for every Origin-Destination market and hence the choice for the passenger is a combination of path and a fare class for the O-D market of interest.

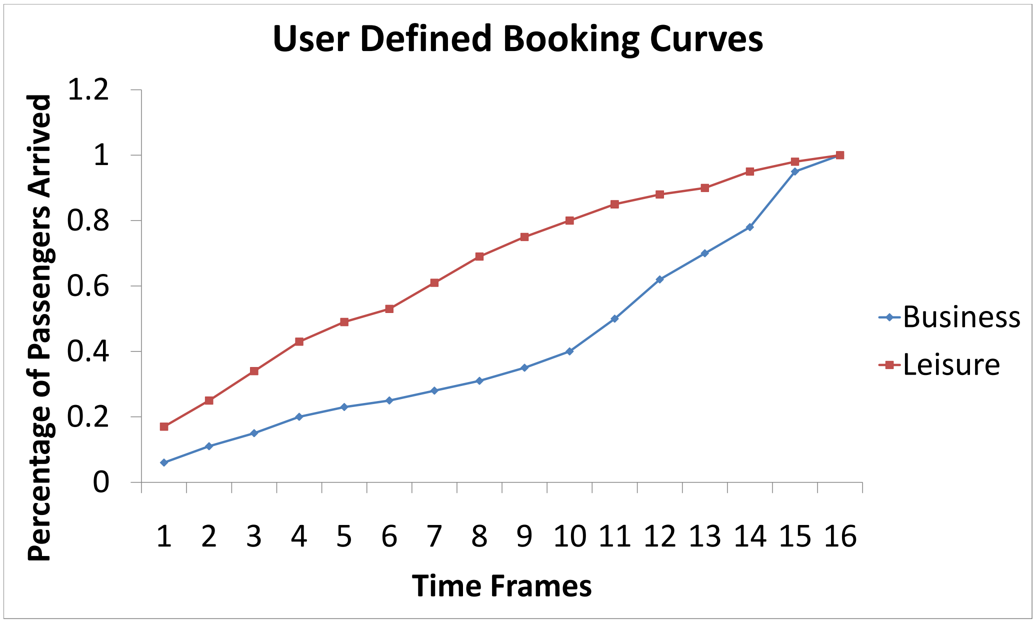

Figure 4-2: PODS Architecture (Source: The PODS Primer)

in four steps – Demand Generation, allocation of Passenger Characteristics, identification of Passenger Choice Set, and modeling the passenger decision. We provide a brief overview of each of these stages, and the reader may refer to Carrier (2003) for more details. The Demand Generation stage encompasses generating average daily air travel demand for every OD pair in the network, based on mean inputs of the "business" demand and "leisure" demand. The user is required to define a booking curve which defines the arrival pattern across various time frames in a booking process as shown in Figure 4-2. The daily demand curve for each passenger type is generated by adding random variability around the the average demand generated.

In the next step, three sets of characteristics are assigned to each passenger: a decision window[4], a maximum willingness-to-pay and a set of disutility costs. The decision window is smaller for the business passengers and larger for the leisure passengers. In the same way, leisure travelers are modeled as highly price sensitive as compared to the business passengers. Sample maximum WTP curves are shown in

[4]The decision window represents a time window within which the traveler is willing to travel. The lower bound of the window is characterized by the earliest departure time the traveler is willing to consider and the upper bound of the decision window is the latest arrival time the traveler will consider.

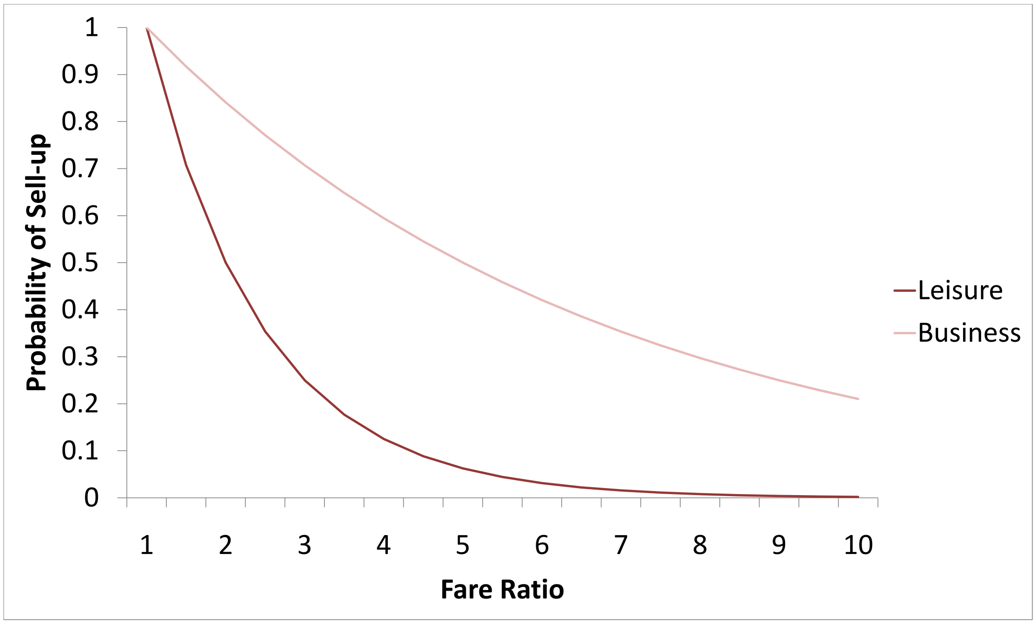

Figure 4-3: PODS Architecture (Source: The PODS Primer)

Figure 4-3.

Any path/class whose fare exceeds the maximum WTP is excluded from that passengers' choice set. Disutility costs are randomly generated from a Gaussian distribution for each passenger type and reflect the sensitivity of the passenger towards schedule, path quality and the fare product restrictions. After the passenger has been assigned a set of characteristics, the passenger is presented with a set of paths and fare classes from which to choose. Some of the options will immediately be removed from the passenger's choice set because of one of the following three reasons: (1) RM system has closed down a path/fare class in the desired OD market, (2) Advance Purchase restrictions cannot be met, or (3) The fare is higher than the passenger's maximum WTP. In the last step, the passenger chooses the option with the lowest generalized total cost (which is sum of the fare and other disutility costs associated with the chosen path/fare class). This information is fed to the airline's RM system to complete the loop as shown in Figure 4-1.

4.2.2 Revenue Management System

We present a higher level overview of the RM systems, as replicated in PODS, in this section and will present specific details in the coming sections about the various components relevant to our work. As discussed before, the RM system modeled in PODS represents a typical third generation RM system with three components: Historical Bookings Database, a Forecaster, and a Seat Allocation Optimizer. The historical bookings database records the path/class of every booking on each airline in the network. The database requires some default bookings at the start of the simulation but those are replaced by actual observed bookings as the simulation progresses (Recall that the first 200 samples are burnt in each trial). The forecaster utilizes the booking data from the database to provide a forecast of future demand by leg/class or by path/class. The data obtained from the database is biased since it only shows the demand that actually booked on a flight. It excludes the passengers who wished to book but could not/did not book due to the availability of various fare classes. The bookings data is, hence, unconstrained and then the forecaster forecasts the future demand. The seat allocation optimizer takes the forecasts generated as inputs and determines the availability of fare classes on every leg/path. PODS has variety of seat allocation optimizers such as EMSRb, DAVN, Heuristic Bid Price and Probabilistic BidPrice. The reader can refer to Belobaba (1992), Belobaba (2002), Cléaz-Savoyen (2005) for a description of these methods. We will compare DP models' performance against EMSRb and DAVN in the next chapter. In the following two sections, we briefly present the forecasting models and DP based seat allocation models used in this thesis.

4.3 Detruncation and Forecasting

The booking data obtained from the historical bookings database is inherently biased as discussed earlier. If a fare class is closed down by the seat allocation model before the end of the booking horizon, actual bookings are less than the actual demand for that fare class. The process used to find the true demand is called detruncation.

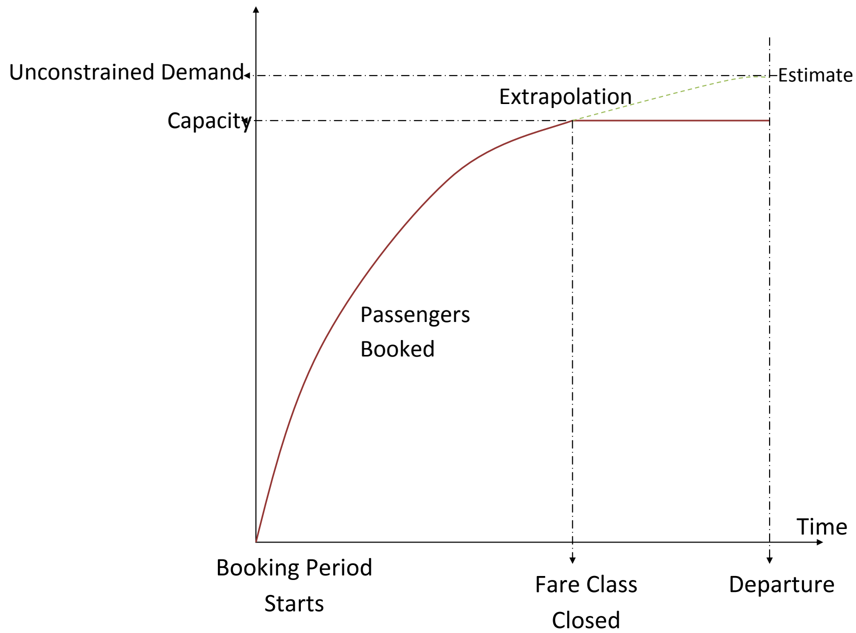

Figure 4-4: Booking Curve Detruncation

There are various approaches to model this problem and reader should refer to the works of Zickus (1998) and Gorin (2000) for a detailed explanation of the various algorithms. The detruncator we will use in this thesis is called booking curve detruncation. This is a relatively straightforward algorithm in which passenger booking curves from previous unconstrained observations are used to project demands of what the forecast would be if there is no capacity constraint. The booking curve detruncation algorithm uses the data from unclosed flights to compute the expected increase in bookings from one period to the next and builds the overall booking curve which can, then, be used to determine the unconstrained demand for closed flights. Figure 4-4 shows the booking curve detruncation algorithm. The reader should refer to Zickus (1998) for mathematical formulation of this algorithm.

4.3.1 Traditional Forecasting Methods

We discussed some common forecasting methods in Section 2.4. In this section, we briefly describe the two forecasting methods which we will use in this thesis and present their mathematical formulations as implemented in PODS. It should be noted that the main purpose of this thesis is not to compare the relative accuracy of these forecasting algorithms; instead we use them to compare the performance of various seat allocation optimizers in a revenue management system.

The standard forecasting method that we will use in this thesis is called pickup forecasting(L'Heureux, 1986). Recall from 2.4 that this method is an elaborate form of time-series forecasting. A time-series forecast would simply be a mean (weighted or unweighted) of final departure bookings on a set of similar flights. Pickup forecasting averages not just the final bookings but also incorporates more information on the number of passengers booked in the intervals before departure. It uses the average number of previous unconstrained bookings and the changes in bookings over time - i.e., the number of passengers that are "picked up" from time period to time period. The demand forecast can be obtained by adding the average pick-up from one time period to another for a fixed number of previous flights to the current bookings(or forecast). A detailed description of this methodology along with a numerical example can be found in Zickus (1998) and Gorin (2000).

The demand can be forecast at either a leg/class level or path/class level. In this thesis, we will refer the former as standard leg based forecasting and the latter as standard path based forecasting.

4.3.2 Hybrid Forecasting

We discussed hybrid forecasting briefly in Section 2.4.2. In this section, we present the tools required to model sell-up and then formulate the equations for Q-forecasting and Hybrid Forecasting.

As explained in Sections 1.3 and 2.4.2, sell-up occurs under a less-restricted fare structure in which passengers tend to buy the lowest available fare. If a passenger is denied booking in his/her choice of fare class, he/she may be willing to pay more for the same flight and "sell-up" to the next higher fare available. We have discussed modeling of maximum WTP in PODS. The seat allocation policy should be designed such that each passenger should pay a fare closest to his/her maximum WTP. Hence, we need to model sell-up in the airline RM system.

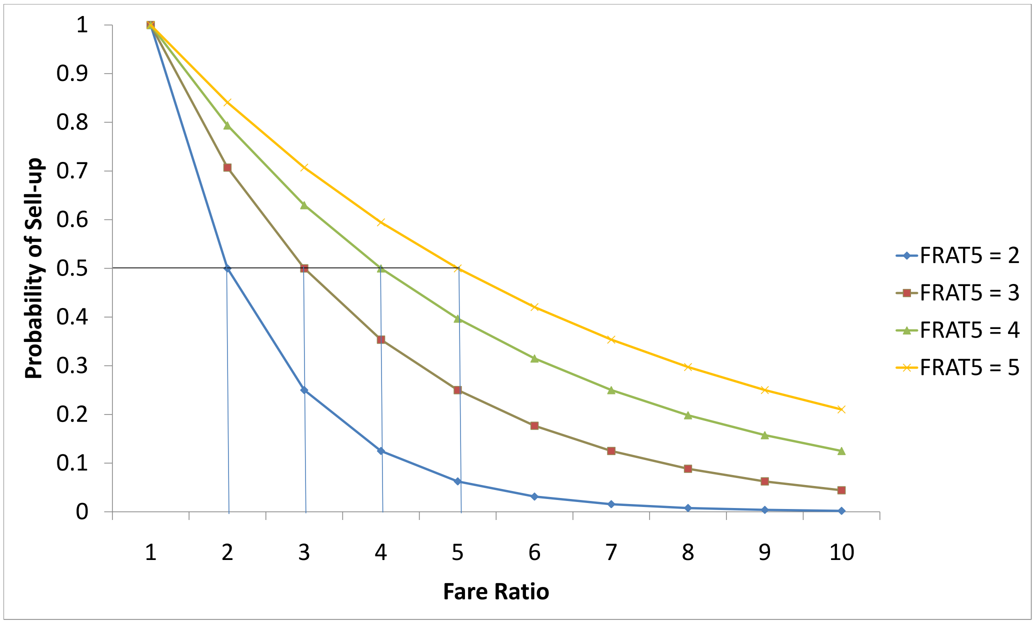

Figure 4-5: FRAT5 Curves

We use the concept of "FRAT5" to model sell-ups in PODS. FRAT5 is defined as the fare ratio of the higher fare to the lowest fare at which 50% of the passengers in lowest class are willing to sell up to the higher fare class.

It is intuitive, based on the definition of FRAT5, that a higher value of FRAT5 represents a higher probability of sell-up to fare classes, as shown in Figure 4-5. We note that this figure is for one time frame and we would expect that passenger's WTP, and hence sell-up probability, increases as we approach the day of departure as price-insensitive business passengers tend to book late in the booking process and they have a higher probability of sell-up. This gives rise to a FRAT5 curve which consists of the FRAT5 values at 16 different time frames and reflects the business/leisure mix at different time frames. A typical FRAT5 curve, FRAT5-"C", that is used in our simulations in this thesis is shown in Figure 4-6.

We note that the traditional methods of forecasting assume independence in fare class demands and do not incorporate sell-up. As described in Section 1.3, traditional methods fail to forecast demand accurately and do not address spiral-down. Belobaba and Hopperstad (2004) develop an approach to forecast the demand in

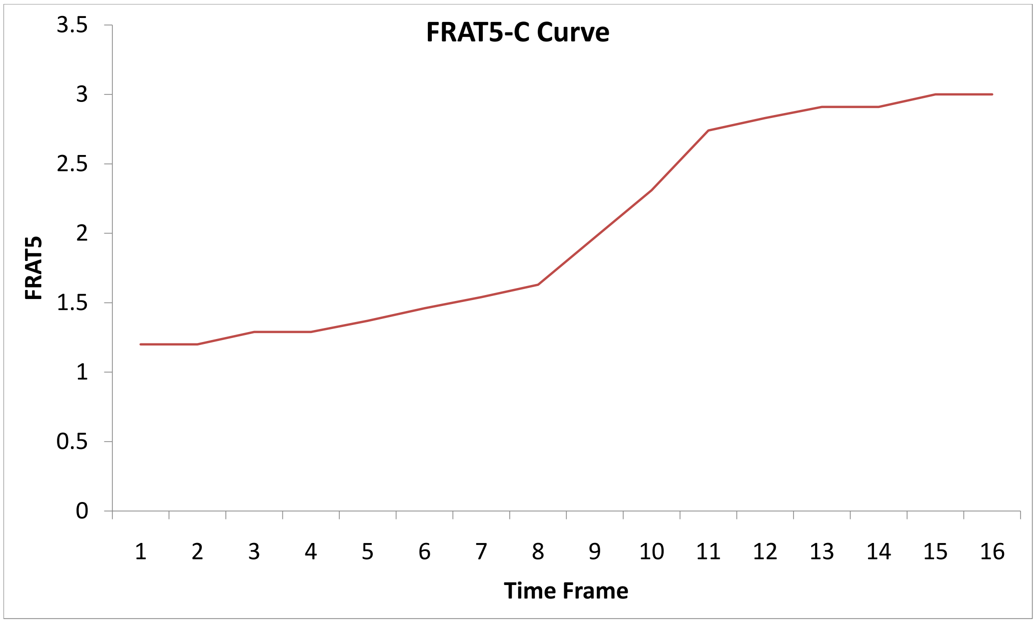

Figure 4-6: FRAT5-C Curve in PODS

unrestricted fare structures. "Q-Forecasting" was introduced in Section 2.4.2 introduces this method. We only present the basic formulation of this algorithm here and the reader is referred to Belobaba and Hopperstad (2004), Reyes (2005) and Cléaz-Savoyen (2005) for detailed description of this method.

The first step in this algorithm is to convert historical bookings into Q-equivalent bookings based on passengers sell-up probabilities. In PODS, the FRAT5 values for each time frame and the fare of each class determines the probabilities of sell-up. Q-equivalent bookings for each class are determined by dividing historical bookings by the sell-up probabilities, as shown below. Let us define the following notation:

$psup_{Q\rightarrow f,tf}$	Probability of sell-up from lowest class Q to class f in time frame tf
$scon_{tf}$	Sell-up Constant in time frame tf
$fare_f$	Fare of Class f
$FRAT5_{tf}$	Fare Ratio at which 50% of passengers will sell-up from Class Q in time frame tf
$hbk_{Q\rightarrow f,tf}$	Estimated equivalent Q-Bookings for fare class f in time frame tf
$hbk_{f,tf}$	Mean Unconstrained Demand for fare class f in time frame tf

We estimate the Q-equivalent bookings as follows:

$$psup_{Q\to f,tf} = e^{-scon_{tf}\cdot\left(\frac{fare_f}{fare_Q}-1\right)}$$

$$scon_{tf} = \frac{\ln(2)}{FRAT5_{tf} - 1}$$

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

The probabilities of sell-up are not easily observable and hence there has been much interest in modeling these sell up probabilities. The FRAT5 curve used in this thesis is a set of predetermined input FRAT5 values. The PODS consortium has been actively involved in estimation methodologies for sell-up probabilities. Readers should refer to Guo (2008) for various methods of estimating the sell-up probabilities.

As discussed in Section 2.4.2, Hybrid Forecasting makes use of two major concepts: Hybrid Forecasting is based on two underlying concepts: (1)The notion of "price-oriented and product-oriented" demand as developed by Boyd and Kallesen (2004), and (2)'Q' Forecasting for price oriented demand as discussed above. Hybrid forecasting classifies each historical booking as a price-oriented booking or a product oriented booking. In PODS, Q-Forecasting makes use of price-oriented booking to generate forecasts of bookings in each undifferentiated fare class. The product-oriented forecasts for each fare class are generated by using traditional pickup forecasting with the product-oriented booking information. The aggregated booking forecasts are generated by combining these two forecasts. Readers should refer to Reyes (2005) who describes the Hybrid Forecasting methodology in detail along with numerical examples.

4.4 Fare Adjustment

Fiig et al. (2010) develop the concept of Fare Adjustment (FA) to improve the performance of seat allocation optimizers by addressing the issue of co-existence of different fare structures on a single leg; typically one of them is a traditional fare structure while the other is an unrestricted (or semi-restricted) fare structure. We recall from Section 2.2.2 that the DAVN method uses displacement adjusted fares[5] to create virtual buckets which are then optimized using seat allocation optimizers like EMSRb. The fare values are fed into a Linear Program which calculates the displacement costs and that displacement cost is used to compute displacement adjusted fare values. The main principle behind fare adjustment is that instead of feeding the LP with original fare values, we should feed the LP with marginal revenue contributions or "adjusted fare" values. The Marginal Revenue contributions are calculated by subtracting a Price-Elasticity cost (PE Cost) from the original fare. One might think of the PE Cost as the cost which accounts for the risk of buy-down. This means that we would use a new adjusted fare (Fare - PE Cost - Displacement Cost) instead of the original displacement adjusted fare (Fare - Displacement Cost) to create virtual buckets. We briefly present the two methods of Fare Adjustment in this section and the readers should refer to Fiig et al. (2010) and Cléaz-Savoyen (2005) for more details on Fare Adjustment.

The two methods of Fare Adjustment are:(1) Marginal Revenue (MR) Formulation, and (2) Karl Isler (KI) Formulation. The underlying concepts in both the formulations are same but the MR formulation is a continuous formulation whereas KI formulation is discrete, as shown below.

$$AdjustedFare^{MR}_{OD,f} = fare_{OD,f} - \frac{fare_{OD,f} \cdot (FAFRAT5 - 1)}{-\ln(0.50)}$$

where,

$AdjustedFare^{MR}_{OD,f}$	MR Adjusted Marginal Revenue of fare f in path OD
$fare_{OD,f}$	OD Fare for Class f in path OD
$FAFRAT5$	FRAT5 Used for Fare Adjustment

[5] Displacement Adjusted Fare = ODFare - Displacement Cost

Karl Isler's discrete formulation is as follows:

$$AdjustedFare^{KI}_{OD,f} = \frac{psup_{Q \to f,tf} \cdot fare_{OD,f} - psup_{Q \to f-1,tf} \cdot fare_{OD,f-1}}{psup_{Q \to f,tf} - psup_{Q \to f-1,tf}}$$

where,

$AdjustedFare^{KI}_{OD,f}$	KI Adjusted Marginal Revenue of fare f in path OD
$psup_{Q \to f,tf}$	Probability of sell-up from lowest class Q to class f in time frame tf
$fare_{OD,f}$	OD Fare for Class f in path OD

Cléaz-Savoyen (2005) argues that the FRAT5 values used for FA should be less than those used for Q-Forecasting by means of a numerical example. An appropriate scaling factor that best describes the passengers WTP to be used in the FA method is used in PODS to model this difference as shown below.

$$FAFRAT5_{tf} = 1 + f5scl \cdot (FRAT5_{tf} - 1)$$

where,

$FAFRAT5_{tf}$	FRAT5 value used for FA in time frame tf
$FRAT5_{tf}$	FRAT5 value used in Q-forecasting in time frame tf
$f5scl$	Scaling factor in time frame tf, $0 \leq f5scl \leq 1$

Figure 4-7 shows the effect of various scaling factors on a sample FRAT5 curve (FRAT5-C). We observe that a full scaling factor of 1.0 is the most aggressive FRAT5 curve and scaling down FRAT5 curve causes it to become less aggressive. A higher scaling factor means that the Marginal Revenue values being fed into the Linear Program to calculate the displacement costs are lower than they would have been with a smaller scaling factor, thereby, increasing the probability that the fare class will get closed down.

In this thesis, we will use the discrete KI formulation along with two different scaling factors of 0.50 and 1.00 to demonstrate the effect of FA aggressiveness on

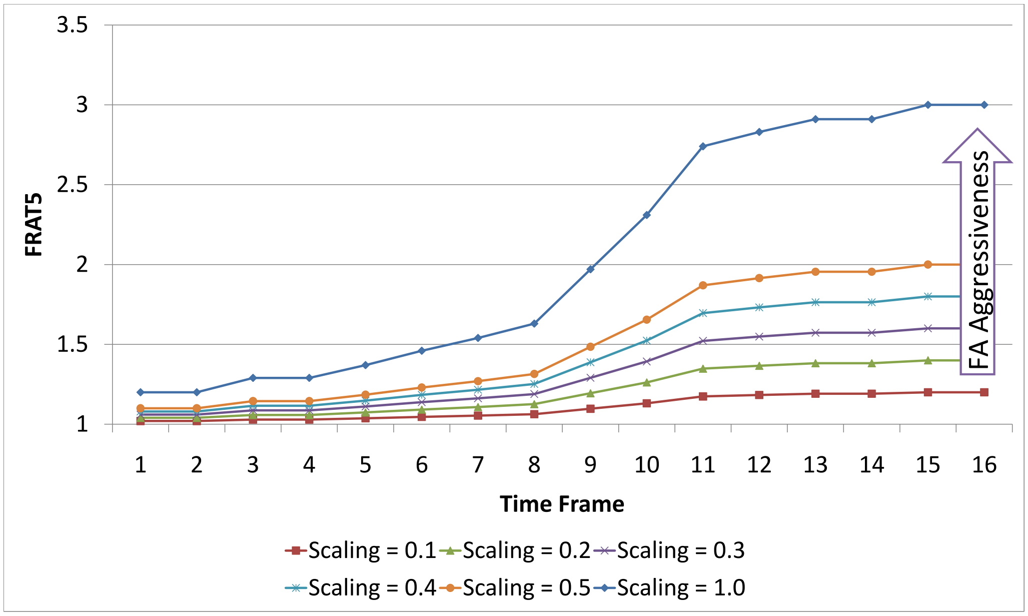

Figure 4-7: FA FRAT5s for Various Scaling Factors

various optimizers.

4.5 Seat Allocation Models in PODS

We discuss some of the seat allocation methods that we will use in this thesis. The first two methods that we use are the traditional RM methods.

4.5.1 EMSRb

As discussed in Section 3.4.2, EMSRb was developed by Belobaba in 1992. In this thesis, we feed EMSRb with traditional pickup forecasts and use booking curve detruncation to determine nested booking limits at the leg level. The optimization model assumes that the demand for each fare class is independent and normally distributed. Figure 4-8 shows the EMSRb optimization process.

As we have mentioned before, EMSRb is widely used as the preferred method of seat allocation on leg level in many airline RM systems. We will use this method as the base case in all our simulations and the performance of other optimizers will be

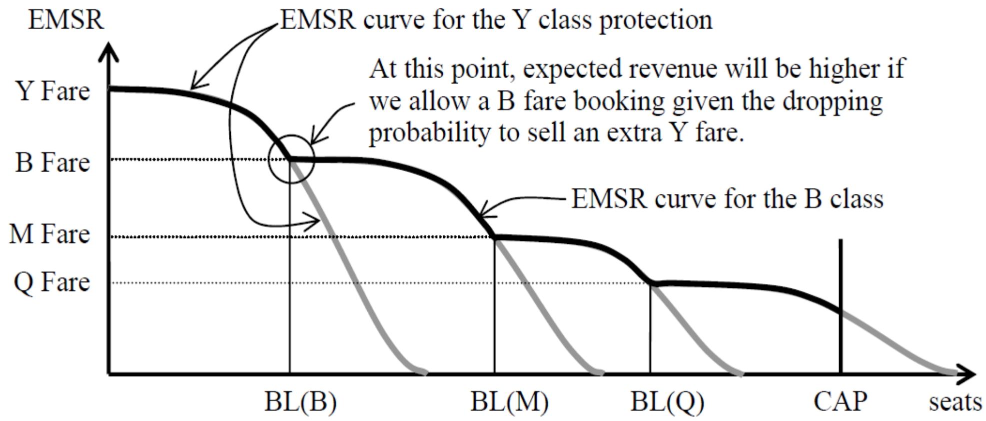

Figure 4-8: EMSR Curve and Booking Limits (Source: PODS Primer)

benchmarked against EMSRb.

4.5.2 DAVN

We introduced Displacement Adjusted Virtual Nesting (DAVN) in Section 2.2.2. For a detailed description of DAVN, readers are referred to Williamson (1992). DAVN is a method of O-D control and hence, it requires forecasts at a path/class level. Belobaba (2002) argues that conventional RM methods do not address two issues which become important when an airline wishes to maximize revenues over a large network: (1) Airline could benefit from accepting a low-yield connecting passenger with a higher network revenue contribution that a higher yield local passenger with smaller network revenue contribution, and (2) Airline could benefit by selling a scarce seat to two local passengers instead of selling the seat to a high-fare connecting passenger, even within the same fare class. DAVN aims to overcome these limitations of leg-based RM methods.

DAVN decreases the fare of the connecting passengers by the amount of network revenue displacement associated with other legs in the path that constitute the flight itinerary. The displacement costs are calculated by solving a deterministic Linear Program (LP) and finding the shadow prices of the capacity constraints. The formulation of the LP is as follows:

$$Max\left(\sum_i \sum_j p_i^j x_i^j\right)$$

subject to:

$$\begin{array}{ll} x_i^j \leq f_i^j & \forall i, j \\ \sum_i \sum_j x_i^j \delta_i^k \leq C_k & \forall k \end{array}$$

where,

p_i^j	Fare for class j on path i
x_i^j	Number of passengers in class j on path i
f_i^j	Forecast for class j on path i
C_k	Capacity on leg k
δ_i^k	Set to 1 if leg k is part of path i and 0 otherwise

For each leg k, we can obtain the displacement cost that corresponds to the marginal revenue of an extra seat on the leg. These values are subtracted from the connecting fares. EMSRb is applied to virtual buckets, formed by ordering fares according to their new value, on each leg to determine the booking limits by bucket. A passenger request for a connecting booking is accepted only if on any leg of the itinerary the corresponding fare is in a virtual bucket that is still available.

For example, consider a flight leg with various fare inputs as shown in Table 4.1. The displacement costs are calculated using the LP formulated above. The displacement adjusted fares are calculated by subtracting the displacement costs from the connect fares.

Class	Local Fare	Connect Fare	Disp. Costs	Disp. Adjusted Fare
1	400	600	150	450
2	300	500	150	350
3	250	400	150	250
4	200	300	150	150
5	150	200	150	50
6	100	150	150	0

Table 4.1: Network Displacement Example (All values in USD)

Table 4.2 shows an example of mapping these fares into 8 virtual buckets. Note that under the scheme used, the class 2 of local fare is ranked above the class 3 of connect fare despite the higher total fare value of class 3 connect passenger.

Virtual Class	Revenue Range ($)	Mapping
1	450+	Connect 1
2	350 - 449	Local 1, Connect 2
3	300 - 349	Local 2
4	250 - 299	Local 3, Connect 3
5	200 - 249	Local 4
6	150 - 199	Local 5, Connect 4
7	100 - 149	Local 6
8	0 - 99	Connect 5, Connect 6

Table 4.2: Virtual Class Mapping

4.5.3 Lautenbacher DP

Lautenbacher DP (LDP) is one of the leg-based DP methods that we will test using PODS in this thesis. The mathematical modeling of this leg-based RM technique has already been discussed in Section 3.4.1. The inputs other than the probabilities of fare class booking requests used in the LDP are known to us (for example, the capacity on the leg, number of time frames, arrival intensity etc) through the problem definition. The probability of a fare class arrival in a certain time frame is computed using the leg/class forecasts of bookings-to-come for a given time frame. We mention one subtle point of implementing this in PODS – if the airline operates in a network and uses DP on a leg basis, the forecaster uses observed bookings to produce the demand forecasts by path which need to be aggregated to leg-based forecasts which LDP can use. Thus, even in a network, LDP is applied at the flight leg level. The reader is referred to Vanhaverbeke (2006) and Tam (2008) for further details on this methodology and some numerical examples.

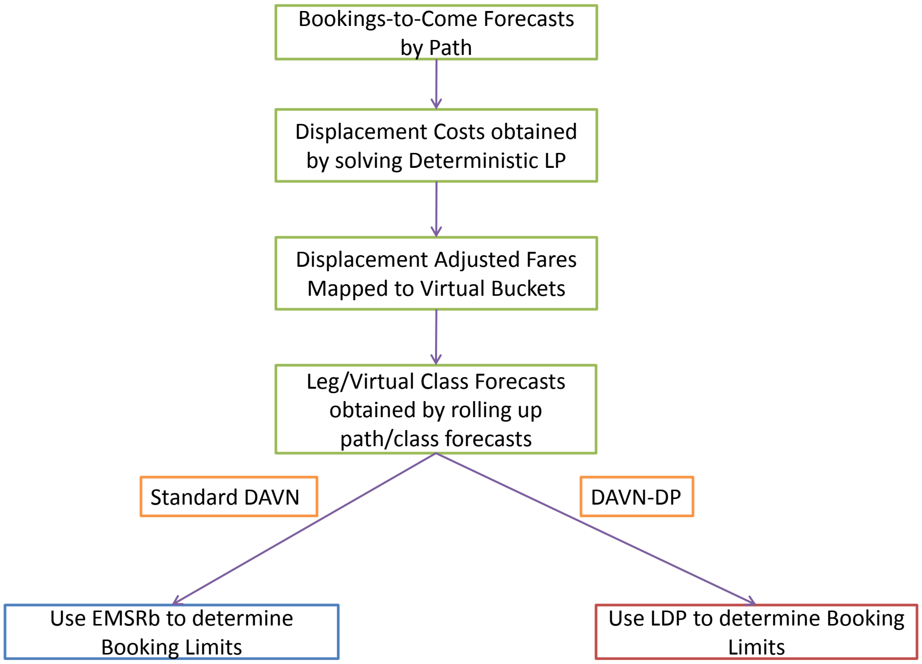

Figure 4-9: Comparison of Standard DAVN and DAVN-DP Process Flow

4.5.4 DAVN-DP

The above leg based DP models have been studied by Vanhaverbeke (2006) and Tam (2008). We introduce a network RM technique using DP, namely DAVN-DP. The idea of using network displacement to adjust the connecting fares was discussed in previous sections. The standard DAVN optimizer, as discussed in Section 4.5.2, utilizes the shadow prices from the deterministic LP program to adjust the connecting fares for displacement costs and then maps the fares into virtual buckets as discussed before. Once the virtual mapping is done, the standard method uses EMSRb at leg level to determine booking limits. DAVN-DP utilizes all the steps of standard DAVN except the last step – it replaces EMSRb with a standard leg based DP (LDP) to determine the booking limits as shown in Figure 4-9. We expect that this optimizer will perform better than the standard LDP, which is a leg based DP, in a network RM environment since we are incorporating displacement costs.

A DP based RM optimizer requires a forecast of bookings-to-come in future time frames instead of total bookings-to-come by fare class. We use a small example

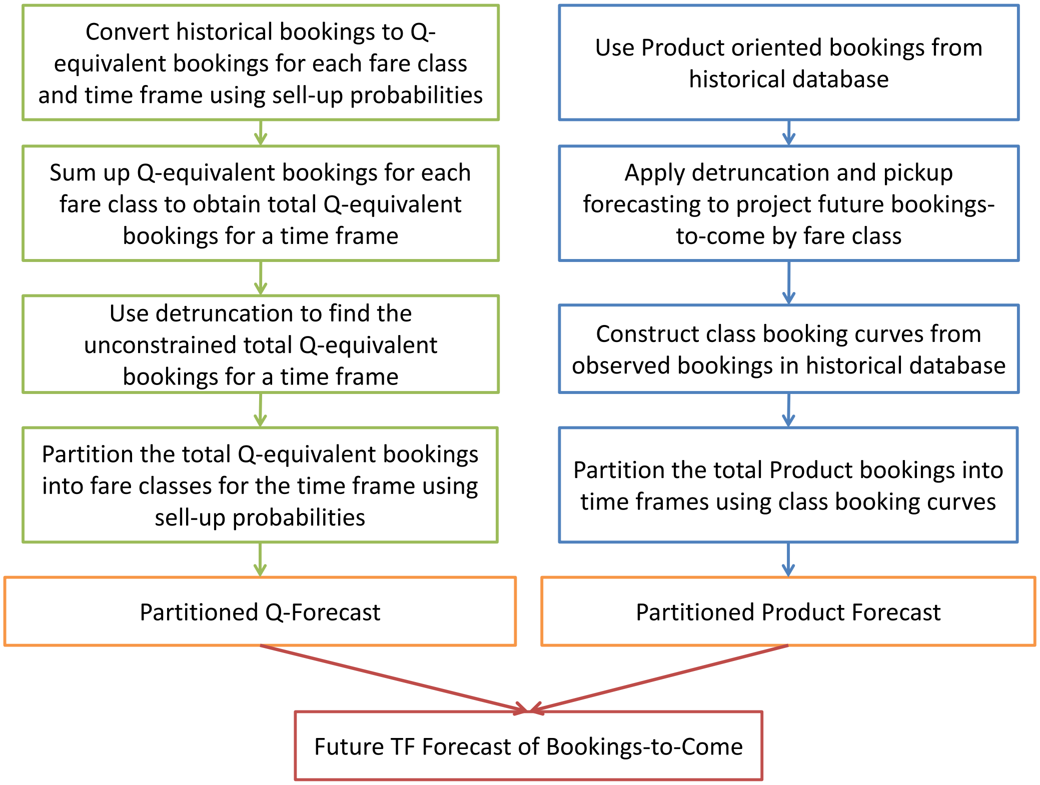

Figure 4-10: DAVN-DP: Forecasting Methodology

to explain the method of constructing future time frame forecasts. We argue that standard forecasting is a special case of hybrid forecasting without the Q-forecast component and hence, a general example of hybrid forecasting would be good to describe the methodology. The steps involved in forecasting are shown in Figure 4-10.

Let us assume that there are four classes (Y, B, M and Q) and four time frames with FRAT5 equal to 1.5, 2, 3 and 4 respectively. Table 4.3 shows the mean unconstrained demands for all fare classes across the four time frames. Based on this data, we can convert it into Q-equivalent bookings for each fare class using the following:

We estimate the Q-equivalent bookings as follows:

$$psup_{Q \to f,tf} = e^{-scon_{tf} \cdot \left(\frac{fare_f}{fare_Q} - 1\right)}$$

$$scon_{tf} = \frac{\ln(2)}{FRAT5_{tf} - 1}$$

Class	Fare	TF1	TF2	TF3	TF4
Y	400	0	0	3	6
B	300	0	2	3	4
M	200	5	4	2	0
Q	100	10	8	5	0

Table 4.3: Mean Unconstrained Demand

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

where:

$psup_{Q \to f,tf}$	Probability of sell-up from lowest class Q to class f in time frame tf
$scon_{tf}$	Sell-up Constant in time frame tf
$fare_f$	Fare of Class f
$FRAT5_tf$	Fare Ratio at which 50% of passengers will sell-up from Class Q in time frame tf
$hbk_{Q \to f,tf}$	Estimated equivalent Q-Bookings for fare class f in time frame tf
$hbk_{f,tf}$	Mean Unconstrained Demand for fare class f in time frame tf

Table 4.4 shows the sell-up probabilities and the Q-equivalent bookings for each fare class in each time frame.

		Sell-up Probability				Q-Equivalent Booking			
Class	Fare	TF1	TF2	TF3	TF4	TF1	TF2	TF3	TF4
Y	400	0.016	0.125	0.354	0.500	0.0	0.0	8.5	12.0
B	300	0.063	0.250	0.500	0.630	0.0	8.0	6.0	6.3
M	200	0.250	0.500	0.707	0.794	20.0	8.0	2.8	0.0
Q	100	1.000	1.000	1.000	1.000	10.0	8.0	5.0	0.0
Total Q- Equivalent Bookings						**30.0**	**24.0**	**22.3**	**18.3**

Table 4.4: Sell-up Probabilities and Q-Equivalent Demand

We use pick-up forecasting and detruncation to estimate total unconstrained Q-equivalent bookings for each time frame and the total Q-equivalent bookings are then partitioned back into separate fare classes by estimating passengers that will sell-up to fare class f but not to $f-1$. We determine the class forecasts by subtracting the potential demand of fare class $f-1$ from the potential demand of fare class f. Let us assume, for the sake of this example, that the detruncated Q-equivalent bookings for the four time frames are 31.25, 25.25, 23.56, and 19.60 respectively. Table 4.5 shows the partitioned forecast which are calculated as follows:

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

Class	TF1	TF2	TF3	TF4
Y	0.5	3.2	8.3	9.8
B	1.5	3.2	3.5	2.5
M	5.9	6.3	4.9	3.2
Q	23.4	12.6	6.9	4.0

Table 4.5: Partitioned Q Forecast

Table 4.5 gives us the partitioned Q forecasts. Partitioned product forecasts are obtained from using the pickup forecasting on the observed product bookings. The PODS consortium is actively doing research on various techniques for characterization of bookings as product oriented booking or price oriented booking. The purpose of this example is to demonstrate the way forecast is computed for DAVN-DP and not to explore various methods to characterize a booking as product/price oriented. Hence, we assume that pick-up forecasting generates the partitioned product forecasts as shown in Table 4.6. The total forecasts for future time frames are obtained by adding the price forecasts (partitioned Q forecasts) and the product forecasts as shown in Table 4.6.

	Q-Eq. Forecast				Product Forecast				Total Forecast			
Class	TF1	TF2	TF3	TF4	TF1	TF2	TF3	TF4	TF1	TF2	TF3	TF4
Y	0.5	3.2	8.3	9.8	0	0	5	5	0.5	3.2	13.3	14.8
B	1.5	3.2	3.5	2.5	0	0	5	5	1.5	3.2	8.5	7.5
M	5.9	6.3	4.9	3.2	0	0	0	0	5.9	6.3	4.9	3.2
Q	23.4	12.6	6.9	4.0	0	0	0	0	23.4	12.6	6.9	4.0

Table 4.6: Future Time Frame Forecast Construction

The results in Chapter 5 have been obtained using the forecasting method described through the numerical example shown above. An alternative way to construct future time frame forecast is to sum up all the partitioned Q forecast by class and add the product future bookings-to-come (BTC) to it to produce a joint BTC forecast. This joint BTC forecast can, then, be partitioned into time frame forecasts using the observed class booking curves.

4.5.5 Unbucketed DP

The number of virtual buckets on each flight leg are fixed (eight) in DAVN-DP and this slightly modified version of DAVN-DP relaxes the hard constraint of 8 path/class buckets. Unbucketed DP (UDP) is a slightly modified version of DAVN-DP because it has 66 path/class buckets on each leg instead of 8 buckets on each leg in DAVN-DP and we expect that UDP will address the issue of sensitivity of DAVN-DP to the number of buckets. Another major difference between DAVN-DP and UDP is in the treatment of adjusted fares. In the backward recursion of the DP, the membership of 8 path/class buckets is static across time frames since the virtual class mapping is done based on current time frame adjusted fares. In theory, we argue that using future time frame adjusted fares should produce optimal revenues and hence, the virtual class mapping should be dynamic across time frames. UDP addresses this issue as illustrated by a small example below.

Numerical Example

Step 0: The Data

- 4 Classes(Y, B, M, and Q)
  - Fares = 400, 300, 200 and 100
- 2 Time Frames
  - FRAT5 = 2.0, 3.0
  - Q Price Demand = 10, 10
  - Y Product Demand = 5, 5

– B Product Demand = 5, 5

Step 1: Obtain the sell-up probabilities

As shown before, we can use the FRAT5 curves and the fare ratios to estimate the sell-up probabilities. The sell-up probabilities for this example are shown in Table 4.7.

Step 2: Obtain the partitioned Q-forecast

We have shown before the method of construction of partitioned Q forecasts by fare class for each time frame using the sell-up probabilities. The partitioned Q forecasts are shown in Table 4.7.

		Sell-up Probability		**Partitioned Q FC**	
Class	QFC	TF1	TF2	TF1	TF2
Y	10	0.125	0.354	1.25	3.536
B	10	0.25	0.5	1.25	1.464
M	10	0.5	0.707	2.5	2.071
Q	10	1	1	5	2.929

Table 4.7: Sell-up Probabilities and Partitioned Q Forecasts

Step 3: Obtain the joint forecast

As seen before, the joint forecasts can be obtained by adding the the partitioned Q forecasts and the product forecasts. The joined forecasts are shown in Table 4.8.

		P_{sup}		Q_{part}		Product FC		**Joint FC**	
Class	QFC	TF1	TF2	TF1	TF2	TF1	TF2	TF1	TF2
Y	10	0.125	0.354	1.25	3.536	5	5	6.25	8.536
B	10	0.25	0.5	1.25	1.464	5	5	6.25	6.464
M	10	0.5	0.707	2.5	2.071	0	0	2.5	2.071
Q	10	1	1	5	2.929	0	0	5	2.929

Table 4.8: Joint Forecasts

Step 4: Obtain the Adjusted Q-Fares and Adjusted Q-Fare Weights

We discussed the KI methodology for fare adjustment before. Using the sell-up probabilities and the fares, we can find the Q-adjusted fares using KI fare adjustment methodology as follows:

$$AdjustedFare_{OD,f}^{KI} = \frac{psup_{Q\to f,tf} \cdot fare_{OD,f} - psup_{Q\to f-1,tf} \cdot fare_{OD,f-1}}{psup_{Q\to f,tf} - psup_{Q\to f-1,tf}}$$

where,

$AdjustedFare_{OD,f}^{KI}$	KI Adjusted Marginal Revenue of fare f in path OD
$psup_{Q\to f,tf}$	Probability of sell-up from lowest class Q to class f in time frame tf
$fare_{OD,f}$	OD Fare for Class f in path OD

Table 4.9 shows the Q-adjusted fares based on the above formulation. The weights for each class are calculated the fraction of Q-partitioned forecast in the total forecast for each class and time frame. The Q-weights are also shown in Table 4.9.

Qpart		**Prod. FC**		**Joint FC**		**Q-Adj. Fares**		**Weights**	
TF1	TF2	TF1	TF2	TF1	TF2	TF1	TF2	TF1	TF2
1.25	3.54	5	5	6.25	8.54	400	400	0.2	0.414
1.25	1.46	5	5	6.25	6.46	200	57.5	0.2	0.226
2.5	2.07	0	0	2.5	2.07	100	-41.5	1	1
5	2.93	0	0	5	2.93	0	-141.3	1	1

Table 4.9: Adjusted Q Fares and their weights

Step 5: Obtain the Joint Future Time Frame Adjusted Fares

The joint adjusted fare can be found by taking a weighted average of the original fares and the Q-adjusted fares with weights as shown in Table 4.9.

$$JointAdjFare_{f,tf} = w_{f,tf} \cdot QadjFare_{f,tf} + (1 - w_{f,tf}) \cdot Origfare_{f,tf}$$

where,

$JointAdjFare_{f,tf}$	KI Joint Adjusted Fare
$w_{f,tf}$	Weight of Q-Adj. Fare for class f in time frame tf
$QadjFare_{f,tf}$	Q-Adjustedf Fare for class f in time frame tf
$Origfare_{f,tf}$	Original Fare for class f in time frame tf

Table 4.10 shows the computed values of joint adjusted fare and we can observe that the adjusted fares are different in different time frames. This shows that the virtual bucketing based on adjusted fares in future time frames is dynamic since the adjusted fares are different for different time frames.

	Original Fare		Q Adj. Fare		Q-Weights		Joint Adj. Fare	
Class	**TF1**	**TF2**	**TF1**	**TF2**	**TF1**	**TF2**	**TF1**	**TF2**
Y	400	400	400	400	0.2	0.414	400	400
B	300	300	200	57.5	0.2	0.226	280	245.1
M	200	200	100	-41.5	1	1	100	-41.5
Q	100	100	0	-141.3	1	1	0	-141.3

Table 4.10: Current Time Frame Adjusted Fares: UDP Methodology

The above example shows the way UDP constructs future TF adjusted fares which differ from time frame to time frame resulting in a dynamic path/class membership in virtual buckets. We use the same example and contrast the way DAVN-DP constructs current time frame adjusted fares and we show that the adjusted fares , and hence the virtual bucket membership, remain constant in all future time frames during one backward trudge of the DP.

	Fares		Forecasts				
Class	**Actual**	Q_{adj}	Q_{Part}	**Prod.**	**Total**	**Q-Weight**	**Joint Adj. Fare**
Y	400	400	4.786	10	14.786	0.323	400
B	300	200	2.714	10	12.714	0.214	278.6
M	200	100	4.571	0	4.571	1	100
Q	100	0	7.929	0	7.929	1	0

Table 4.11: Current Time Frame Adjusted Fares: DAVN-DP Methodology

Table 4.11 shows the DAVN-DP methodology which utilizes the current time frame adjusted fare and calculates the joint adjusted fare which remains constant throughout the backward trudge of DP and hence the membership of path/class in the virtual buckets is static.

We incur a very sharp runtime penalty by increasing the number of buckets in UDP but we expect that it will help us understand if DAVN-DP is very sensitive to the number of virtual buckets. UDP also presents a theoretically better alternative of using future time frame adjusted fares when using fare adjustment. We will evaluate UDP alongside DAVN-DP in the next chapter when we present the results of PODS simulations.

4.6 Summary

We presented an overview of PODS in this chapter to introduce the simulation environment before we present results in the next chapter. The basic architecture of PODS was discussed and two of its main components – Passenger Choice Model and the RM system were discussed. We also discussed standard (leg and path) and hybrid forecasting along with fare adjustment techniques. Seat allocation optimizers relevant to this thesis like EMSRb, Standard DAVN and LDP were briefly discussed. We introduced two new DP based network RM methods – DAVN-DP and UDP. In the next chapter, we will present results using these optimizers with the two forecasting methods discussed in this chapter.

Chapter 5

PODS Simulation Results

In Chapter 3, we discussed the results obtained in a single leg case with no competition and forecasting effects. We now extend our analysis to a more complex airline network setting. We choose a relatively symmetric network in which airlines offer very similar fare structures and schedules because we want to ensure that the main distinction between the two competing airlines is their RM system so that we can further evaluate the performance of different optimizers discussed in the previous chapter. This chapter presents the results from our PODS simulations to test the performance of Dynamic Programming based approaches in Revenue Management systems. The chapter is divided into three sections. The first section introduces the network used for the simulations performed in this thesis. The second part will present the findings in the network with semi-restricted fare structures and the last section focuses on the results obtained in the network with fully restricted fare structures.

We will compare the performance of DP models described in the previous chapter relative to a "base case". The base case for the study is chosen to reflect the standard RM systems used by airlines. We chose EMSRb with standard leg/class forecasting (Standard EMSRb) as the base case for the evaluation of leg based DP methods and we use DAVN with standard path/class forecasting (Standard DAVN) as the base case for DP based network RM models.

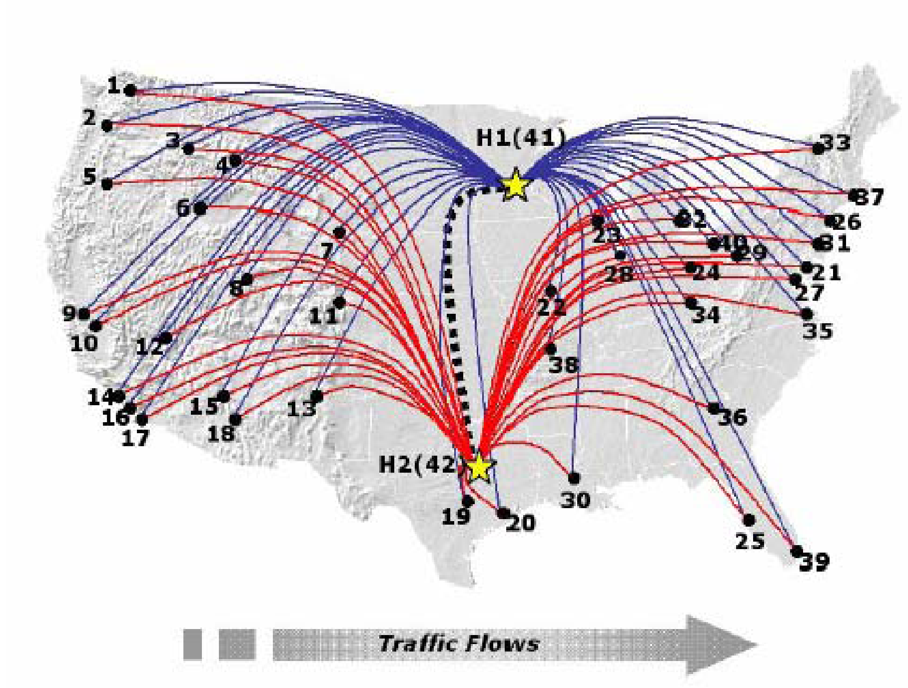

Figure 5-1: Schematic for Network D6

5.1 Network D6 Overview

Network D6 is a simplified representation of a US domestic airline network with two competing hub-and-spoke carriers. Each of the hubs has 20 eastern spoke cities and 20 western spoke cities. Airline 1 is located at the Minneapolis-Saint Paul International Airport (MSP), whereas the competing airline, or Airline 2, is located at the Dallas-Fort Worth International Airline (DFW), as shown in 5-1. As shown in Figure 5-2, the traffic flows from the western cities to the eastern cities with three bank times[1] — 10:30 AM, 2:00 PM, and 5:30 PM. Both the airlines also offer non-stop inter-hub flights. There are 252 flight legs and a total of 482 O-D markets with each airline experiencing head-to-head competition in all of their markets.

Both airlines offer identical fare structures consisting of 6 fare classes which are differentiated by fare price and restrictions. As discussed before, we will present results for two cases – one with both airlines offering semi-restricted fare structures and the other with both airlines offering fully restricted fare structures. The average

[1]Bank time refers to the time when eastbound passengers on different flights arriving from western cities are consolidated at the hub on flights departing for the east coast cities

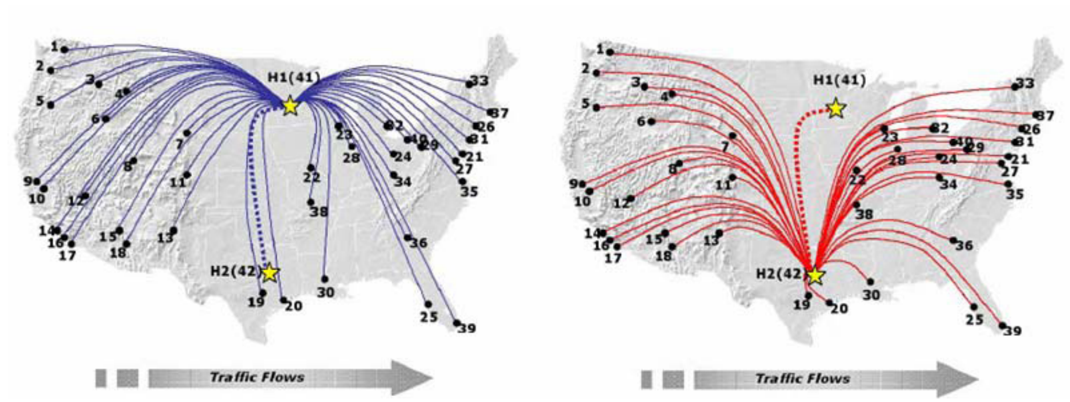

Figure 5-2: Route Map for Airline 1 (left) and Airline 2(right)

fare values range from as low as \$101 for the lowest Class 6 to \$413 for the highest Class 1. We will discuss the restrictions associated with the fare structures in subsequent sections.

5.2 Semi-Restricted Fare Structure

We present the results in a semi-restricted fare structure, as shown in Table 5.1, in network D6. Most common fare structures incorporate the following restrictions to segment the demand — (1) Saturday Night Minimum Stay, (2) Change Fee and (3) Limited Refundability on Canceled Bookings. Apart from these restrictions, another restriction, known as Advance Purchase restriction, is commonly used by the airlines. In the semi-restricted fare structure shown below, the advance purchase restriction is shown by "AP" and the other three restrictions are characterized by R1, R2 and R3. We observe that the restriction R1 corresponding to "'Saturday Night Minimum Stay" is completely removed while the other two restrictions along with advance purchase restriction are present.

5.2.1 First Choice Only Choice

We discuss the results obtained by assuming that the forecasting methodology used by both the airlines is the standard pickup forecasting approach described in Section 4.3.1. We impose an additional constraint in PODS, called First Choice Only

Fare Class	Avg. Fare	AP	R1	R2	R3
1	412.85	0	0	0	0
2	293.34	3	0	1	0
3	179.01	7	0	1	1
4	153.03	14	0	1	1
5	127.05	14	0	1	1
6	101.06	21	0	1	1

Table 5.1: Network D6 Semi-Restricted Fare Structure

Choice (FCOC), which means that a passenger travels only if the first choice requested by the passenger is available and hence the passenger is counted as a spilled passenger (or a No-Go) if the passenger's first choice is not available. A consequence of such a constraint is that there is no passenger sell-up, recapture and spill-in within an airline or across airlines. We believe that FCOC provides us a with a baseline simulation setting to compare DP and EMSRb since it eliminates the effect of competition, inter-airline spill, and sell up.

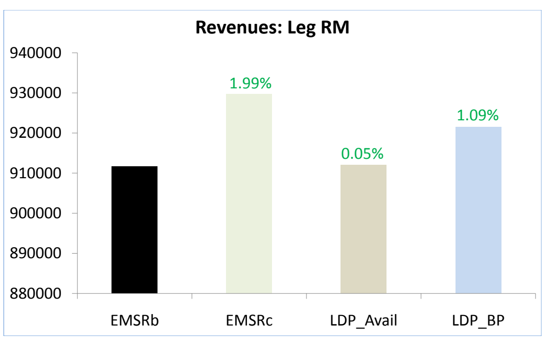

(a) Revenues

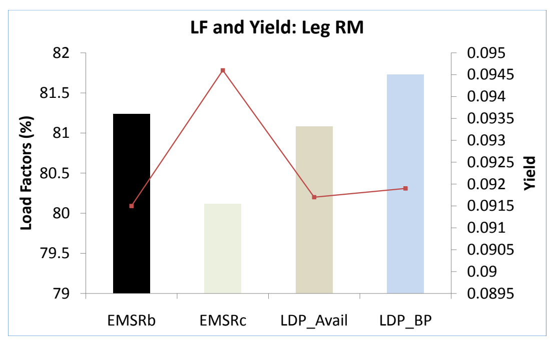

(b) Load Factors

Figure 5-3: Revenues and Load Factors (a) Revenues and (b) LF and Yield

Figure 5-3 shows the results for leg RM systems. "EMSRc" is used to refer the use of bidprice control[2] with EMSRb. We observe that DP outperforms standard EMSRb when used at a leg level but EMSRc (EMSRb with bidprice control) outperforms the leg-based DP irrespective of the control (Bidprice Control or Availability Control) used with DP. DP with Bidprice control outperforms standard EMSRb by

[2]Bidprice Control was described in Section 2.2.2

1.1% whereas DP with availability control increases revenues by 0.05% over EMSRb which is in line with our previous results discussed in Chapter 3. EMSRc has the highest yields and lowest load factors. We also observe that DP with bidprice control has a higher load factor and a lower yield than EMSRb. We look at closure rates for EMSRc, DP with bidprice control and EMSRb to better understand these results.

Figure 5-4 shows the closure rates for EMSRb, EMSRc and DP with bidprice control. We observe that, in later time frames, DP is more open in Class 1 when compared to EMSRb. A higher closure rate for EMSRb in the later time frames translates to fewer Class 1 passengers and contributes to lower yields. In Class 6 closure rates, we observe the same phenomenon as observed in Chapter 3 (in a single flight leg simulation) — DP is more closed in Class 6 than EMSRb in the earlier time frames but DP is more open than EMSRb in later time frames. We also observe that EMSRc is more open in Class 1 and more closed in Class 6 than DP which explains more higher class passengers and less lower class passengers resulting in higher yield for EMSRc when compared to DP. The closure rates explain the reason behind more higher class passengers and less lower class passengers in EMSRc as compared to DP and EMSRb.

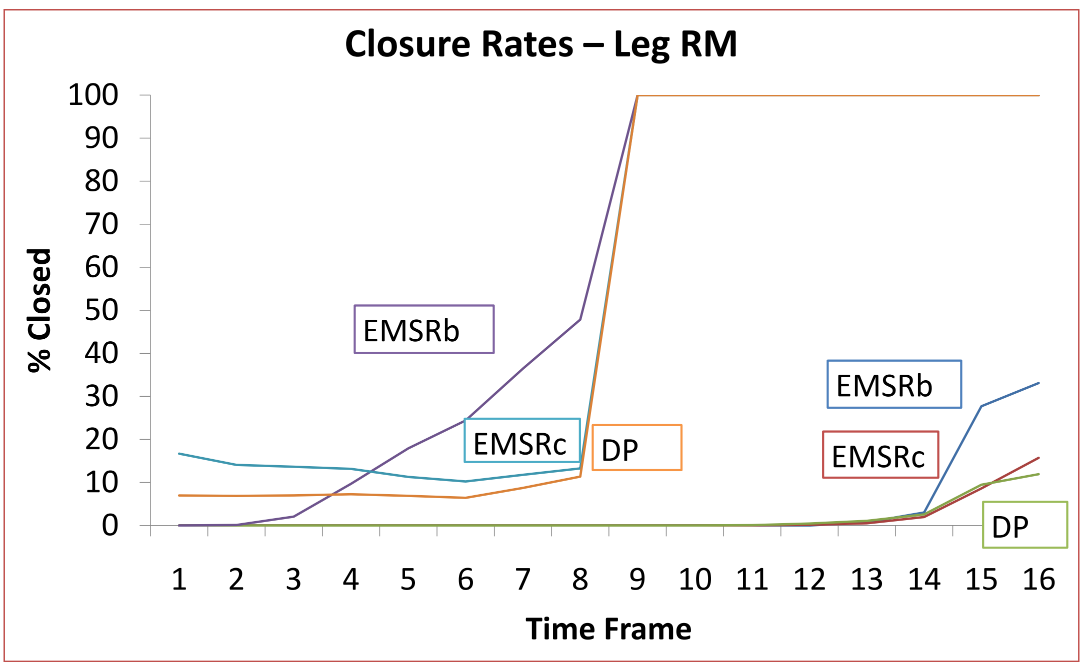

Figure 5-4: Closure Rates for Std. EMSRb and LDP with Bidprice Control

Class	EMSRb	EMSRc	LDP-BP
1	35315	18864	29151
2	9726	4276	13540
3	13252	9970	16344
4	0	0	0
5	14158	10583	13891
6	70729	81353	57490
Total	143180	125046	130416

Table 5.2: No-Go Passenger Revenue (Spill Out)

Table 5.2 shows the spill-out (or No-Go) revenues for EMSRb, EMSRc, and DP with bidprice control. The table shows that overall EMSRc spills out the least passengers (which translates to revenues) when compared to other two optimizers. EMSRc also spills out the least passengers in higher classes despite spilling out more lower class passengers. DP spills out less overall passengers as compared to EMSRb. These results supplement the insights gained before using the closure rates. This explains the higher load factor obtained by DP with bidprice control and EMSRc when compared with standard EMSRb.

Figure 5-5 shows the results of using network RM techniques. The figure shows that the benefit of Airline 1 (AL1) moving to standard DAVN from EMSRb is a revenue increase of 0.67%. The use of DP models in network RM does not generate additional revenues over and above standard DAVN. We observe that DAVN-DP with availability control and UDP underperform relative to standard DAVN while DAVN-DP with bidprice control marginally outperforms standard DAVN. DP methods in the network RM system increase load factors and decrease yields relative to standard DAVN.

We leave the detailed analysis of network RM techniques for the next section where, instead of using FCOC, we use the PODS passenger choice model. The results from the FCOC experiment with standard forecasting are summarized below:

- Leg RM
  - EMSRc (EMSR with bidprice control) performs better than any other

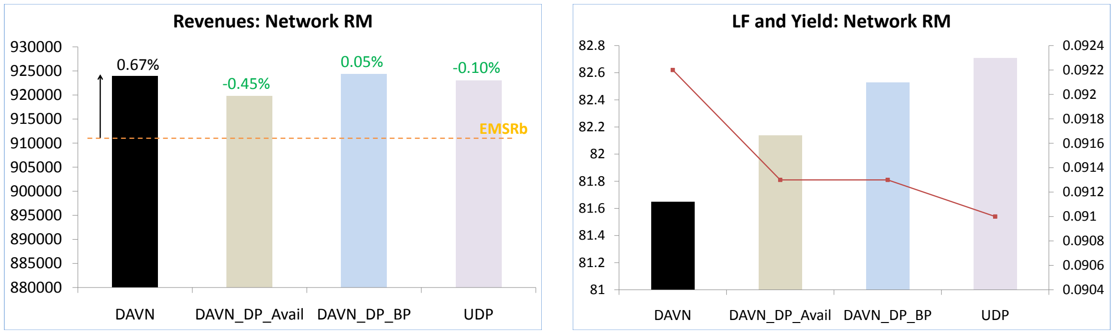

(a) Revenues (b) Load Factors

Figure 5-5: Revenues and Load Factors (a) Revenues and (b) LF and Yield

method tested

- DP performs slightly better than standard EMSRb
  - DP with Bidprice control works better than DP with Availability control
- DP produces higher yields as compared to EMSRb

- Network RM
  - A move from EMSRb to DAVN increases revenues by 0.67%
  - DAVN-DP with bidprice control produces higher revenues than Standard DAVN

- The benefits obtained by DP exclude the effects of sell-up and competition since the simulations assumed First Choice Only Choice

5.2.2 PODS Passenger Choice

In the previous set of experiments, we used a simple choice model – First Choice Only Choice. The results presented in this section use a more elaborate form of passenger choice model as described in the previous chapter. We argue that these experiments will help us capture the impact of competition and passenger sell-up. As before, the base case for leg based RM systems is EMSRb with standard leg/class forecasting and the base case for network RM system is DAVN with standard path/class forecasting.

Competitor Uses Standard EMSRb

The experimental set-up is shown in Table 5.3. In the base case for the leg RM experiments, both Airline 1 and Airline 2 use EMSRb with standard leg/class forecasting. In the base case for the network RM experiments, Airline 1 uses DAVN with standard path/class forecasting while Airline 2 uses EMSRb with leg/class forecasting.

	Leg RM		**Network RM**	
	Airline 1	Airline 2	Airline 1	Airline 2
Base	EMSRb	EMSRb	DAVN	EMSRb
I	EMSRc	EMSRb	DP/Availability Control	EMSRb
II	DP/Availability Control	EMSRb	DP/Bidprice Control	EMSRb
III	DP/Bidprice Control	EMSRb	UDP	EMSRb

Table 5.3: Experimental Set-up

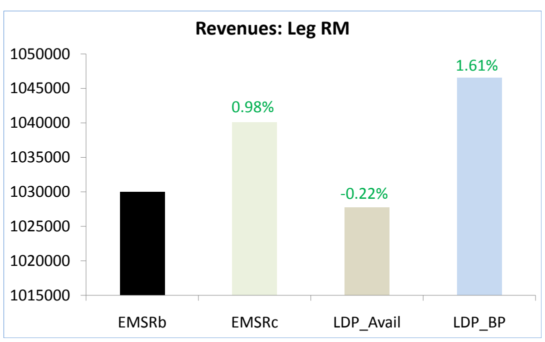

(a) Revenues

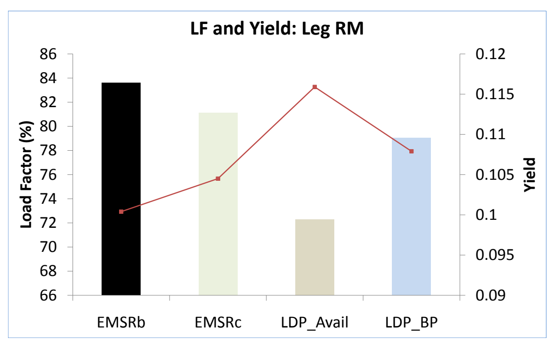

(b) Load Factors

Figure 5-6: Revenues and Load Factors (a) Revenues and (b) LF and Yield

Figure 5-6 shows the revenues and load factors for leg RM systems. We observe that, unlike in FCOC, DP with bidprice control outperforms all the optimizers tested and increases the revenue by 1.6% over EMSRb. We also observe that EMSRb used with bidprice control also outperforms standard EMSRb (EMSRb with availability control) by 1%. It is an important result that we observe here – standard EMSRb outperforms DP with availability control. The load factors obtained by DP and EMSRc are lower than EMSRb but their yields are higher than the EMSRb yield.

We show the fare class mix in Figure 5-7 to explain the results of DP with availability control. We can see that DP with availability control accepts very few Class 6 bookings relative to other optimizers tested and accepts highest number of Class 1 bookings. As seen in all the results in Chapter 3, DP with availability control is too aggressive in protecting seats for high-yield Class 1 passengers and loses Class 6 passengers. This argument is also supported by the load factor graphs that we saw before – DP with availability control has the highest yields and lowest load factors.

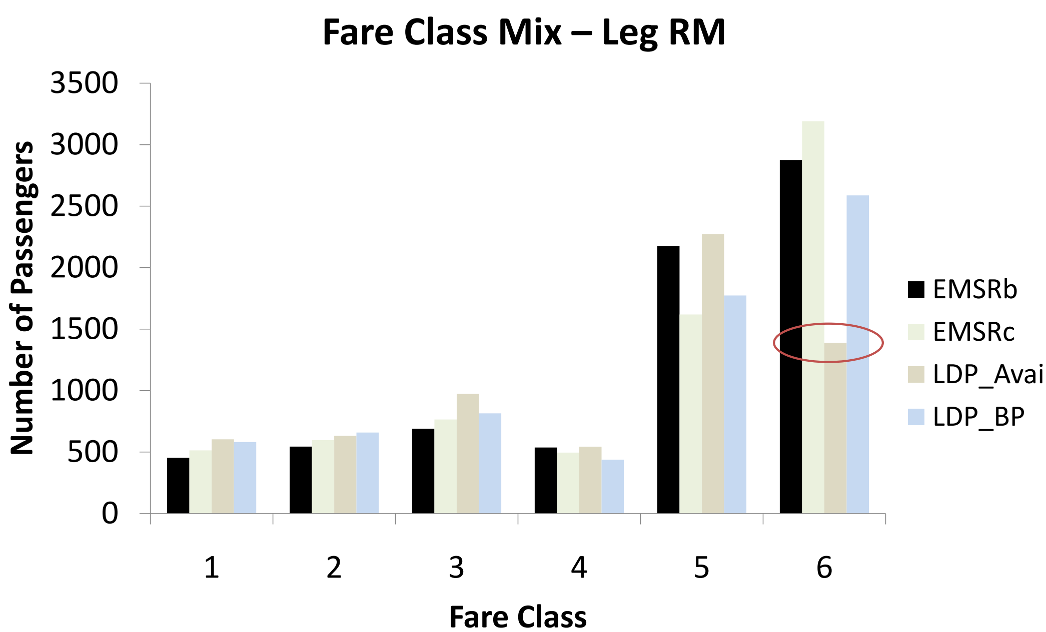

Figure 5-7: Fare Class Mix for Leg RM

Figure 5-8 shows the revenues and load factors for network RM systems. We observe that DP based network RM systems outperform standard DAVN. It can be noted that a move to a network RM system (standard DAVN) increases revenues by about 0.70%. The DP based optimizers outperform standard DAVN by 0.85%–1.70%. We observe that the bidprice control works better than availability control as was the case in leg RM systems. The load factors obtained by DP methods are higher than standard DAVN while the yields are lower compared to standard DAVN.

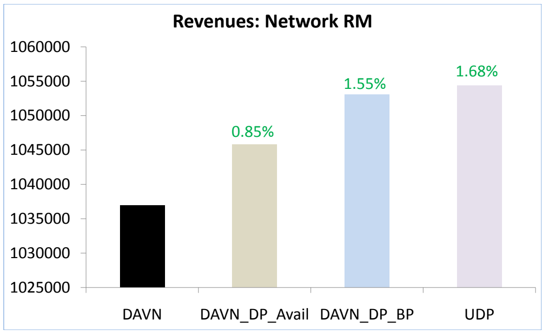

(a) Revenues

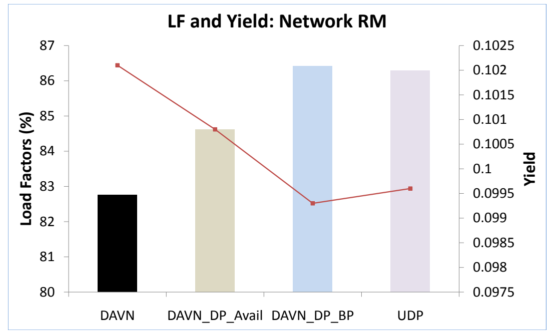

(b) Load Factors

Figure 5-8: Revenues and Load Factors (a) Revenues and (b) LF and Yield

The fare class mix for network RM methods is shown in Figure 5-9. We observe that the DP based methods accept more Class 6 bookings than standard DAVN where as standard DAVN accepts more higher class bookings. This intuitively explains the lower yield but higher load factors for DP methods as compared to DAVN.

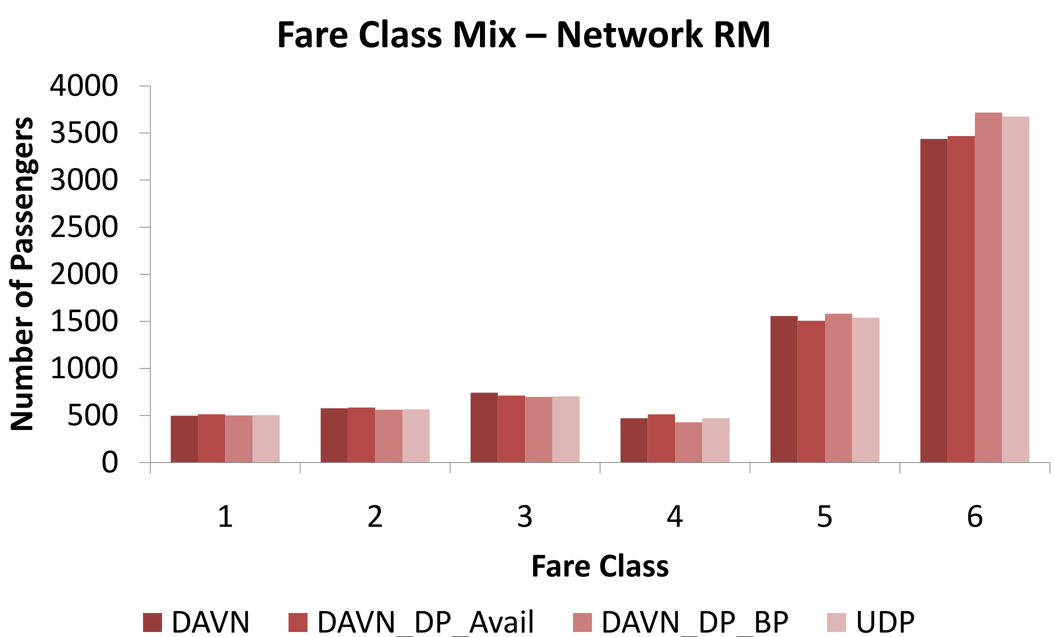

Figure 5-9: Fare Class Mix for Leg RM

Competitor Uses Same RM System as Airline 1

In the previous section, we discussed the results in which competitor (Airline 2) used EMSRb in its RM system in all the test cases. In this section, we present the results in which Airline 2 uses the same optimizer as Airline 1. We believe that such an experiment will give us insights into the effect of competition on the performance of DP. We have discussed the aggressive seat protection policy in DP in Chapter 3 and we saw it in previous experiment as well. This set of experiments will help us in evaluating the effect of competition in the performance of DP methods.

Figure 5-10 shows the revenues, load factors and yields of Airline 1 when the competitor uses the same RM system as Airline 1. The results show that DP outperforms standard EMSRb at leg level. DP with availability control generates the maximum revenues among all the leg based RM methods shown. We observe that when both the airlines use DP with availability control there are big gains in revenues due to symmetry. Both the airlines use an aggressive optimizer in their RM system which leads to a situation of "cooperation" such that both airlines benefit due to aggressive nature of their RM systems. As discussed before, the use of DP with availability control reduces the Class 6 availability in the initial time frames and hence as there are less lower class seats available in the system, passengers are forced to buy fares closer to their willingness-to-pay since now even Airline 2 is more closed in Class 6 by virtue of it using DP. However, we believe that such a cooperation does not exist in the real world but this example shows that by being intelligently aggressive, the airlines can benefit.

Figure 5-11 shows the fare class mix. As seen before, leg based DP methods accept fewer Class 6 bookings leading to lower load factors as compared to standard EMSRb. In Network RM, a move from standard EMSRb to standard DAVN increases revenues by 1% for Airline 1. Although DP methods outperform standard DAVN, the gains by these methods are much less than when Airline 2 used standard EMSRb as discussed in previous section. Figure 5-12 shows these results and we can observe that all the three DP methods outperform standard DAVN but the revenue gains obtained in this

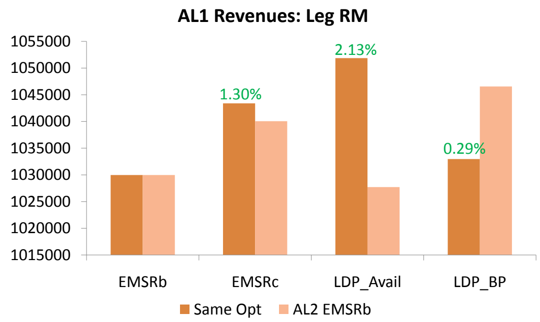

(a) Revenues

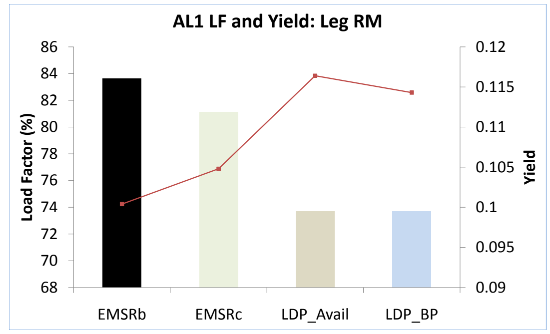

(b) Load Factors

Figure 5-10: Revenues and Load Factors (a) Revenues and (b) LF and Yield

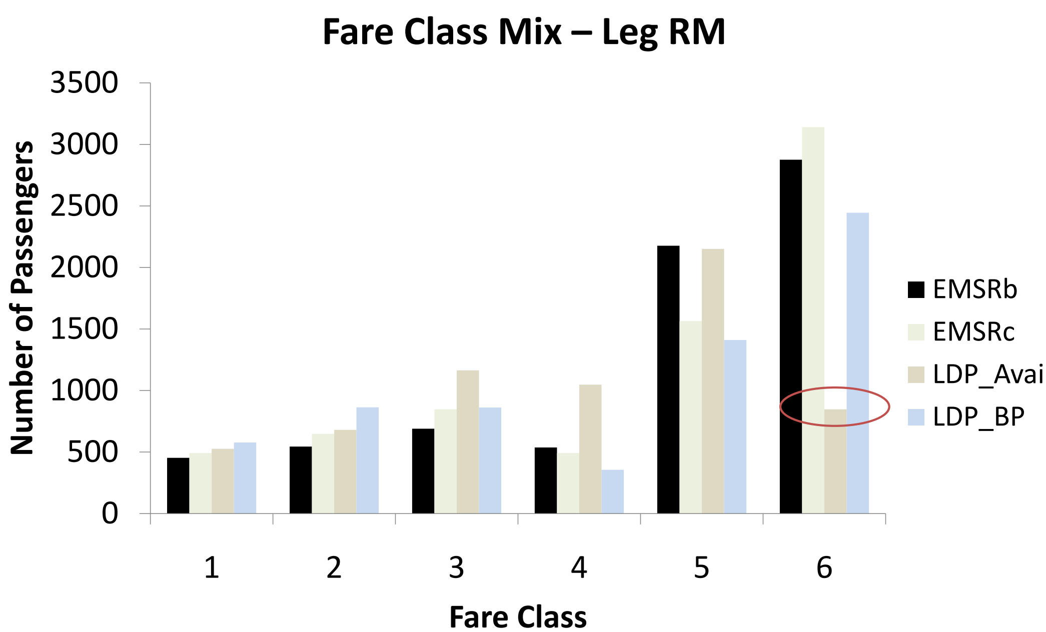

Figure 5-11: Fare Class Mix for Leg RM

case are smaller than the case when airline 2 uses EMSRb.

Discussion of Results

The performance of DP based optimizers was studied in a single flight leg example in Chapter 3. We discussed the results that showed that though DP is a theoretically appealing optimizer, its performance was comparable (if not worse) to EMSRb. We argued, by way of sensitivity analyses, that the benefits depend on various factors

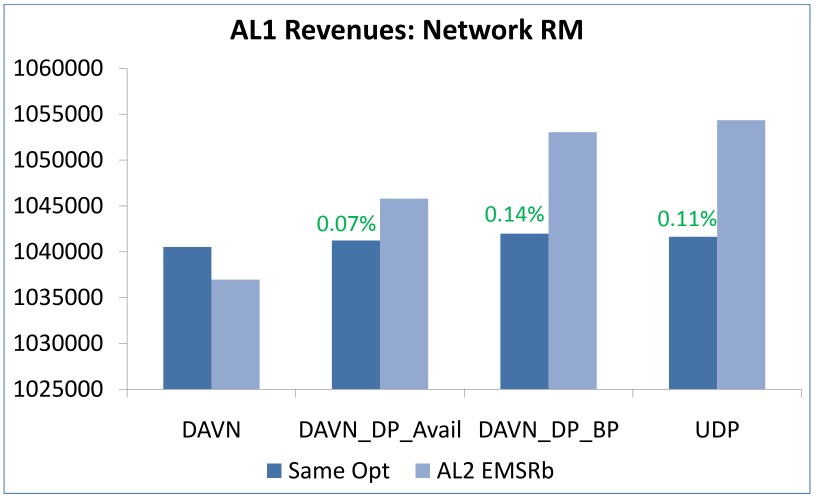

(a) Revenues

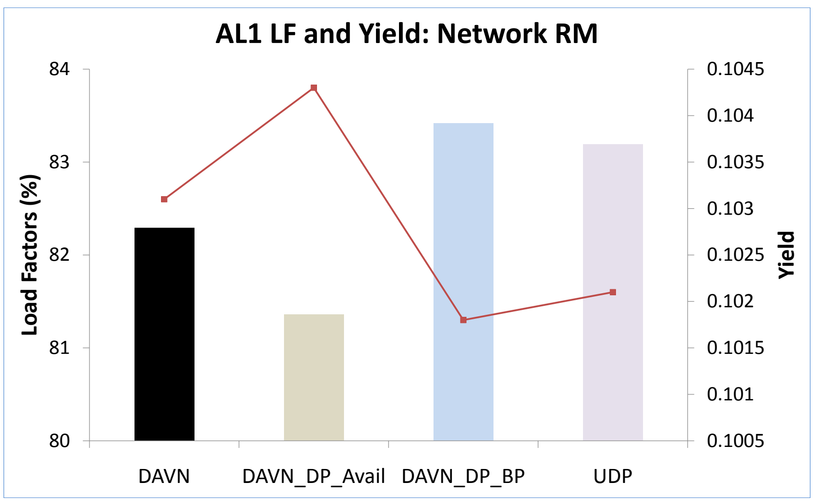

(b) Load Factors

Figure 5-12: Revenues and Load Factors (a) Revenues and (b) LF and Yield

like realized demand variability, fare ratios, capacity relative to mean demand, and the passenger arrival process. We concluded the chapter with the comment that high demand variance adversely affects the performance of DP based methods and the benefits of having a dynamic booking limit policy (as in DP) are negated by the fact that the assumed demand variance in DP models is much less than the variance realized in most airline forecasting systems.

However, in our simulations shown above, we observe that DP performs better than EMSRb in most of the scenarios as shown above. We have discussed before that these experiments would help us understand the effects of competition and passenger sell-up on the performance of DP methods. In the subsequent paragraphs, we explain that though DP based models might not be better optimizers than EMSRb as was shown in Chapter 3, they have inherent properties like aggressive seat protection limits which help them outperform standard EMSRb in a competitive simulation due to factors like competitive feedback and passenger sell-up.

The results discussed above show that the move of Airline 2 to a more sophisticated RM system adversely affected revenues of Airline 1 relative to the case where Airline 2 uses standard EMSRb. To understand the effect of competition, we compare the two sets of results for the network RM systems — first set when Airline 2 uses standard EMSRb and the second set when Airline 2 uses the same RM system. We analyze the results for DAVN-DP with bidprice control as a representative example.

We have seen before that network based DP methods have a higher load factor and lower yields than DAVN because they accept more lower class bookings and fewer higher class bookings. The main intuition behind the competitive analysis done here is that the use of a sophisticated RM approach by Airline 2 affects Airline 1 revenues since Airline 2 will capture more lower class passengers when it moves to DP based optimizer and hence Airline 1 will gain fewer lower class passengers due to less inter-airline spill-in while Airline 2 will gain more lowest Class 6 passengers. The subsequent analysis proves the above intuition based on numerical results obtained through the simulations.

Figure 5-13(b) shows the fare class mix for the two airlines serving the network in the two scenarios as described above. We observe that the total number of passengers in the system remained almost the same when Airline 2 switched from standard EMSRb to DP with bidprice control (recall that Airline 1 also uses DP with bidprice control) and there was just a shift of passengers from one airline to the other. Airline 1 lost about 190 passengers and Airline 2 gained about 190 passengers. We observe that Airline 1 lost most passengers in Classes 5 and 6. Airline 2 lost some passengers from Classes 4 and 5 but gained much more passengers in Class 6.

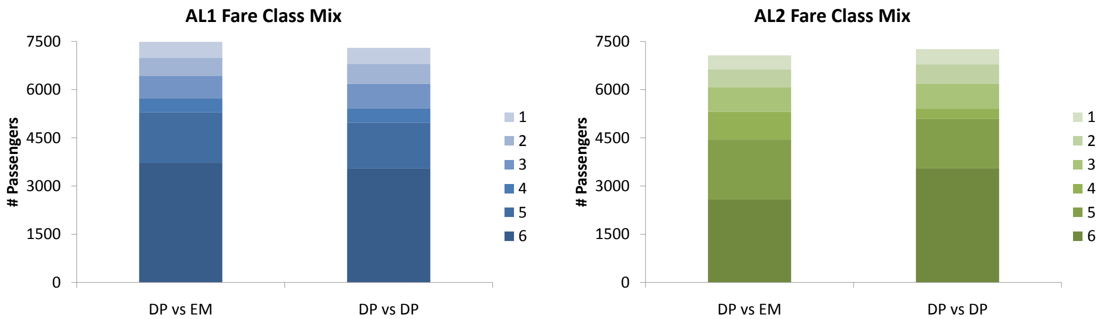

(a) Airline 1 (b) Airline 2

Figure 5-13: Revenues and Load Factors (a) Revenues and (b) LF and Yield

Figure 5-14 shows the breakdown of Airline 2 passengers by fare class in the two scenarios. For classes 4 and 5, we observe that there is a decrease in recaptured passengers by Airline 2 and a decrease in spilled passengers from Airline 1 to Airline 2 when the latter moves from EMSRb to DP. For Class 6, there is an increase in spill-in, recapture, and first choice passengers as Airline 2 switches the RM system from EMSRb to DP with bidprice control. Airline 2 gains most passengers in Class 6 from higher acceptance rate of Class 6 passengers whose first choice was Class 6 on Airline 2 as was discussed before.

Figure 5-15 shows the breakdown of Airline 1 passengers by fare class in the two scenarios as described before. For Class 5, we observe that there is a decrease in recaptured passengers by Airline 1. Airline 1 also gains fewer passengers from Airline 2 due to spill-in as Airline 2 switches to DP from EMSRb. We observe the same phenomenon in Class 6 except that the loss to Airline 1 from fewer spilled-in passengers from Airline 2 is bigger as compared to Class 5.

Figure 5-16(a) shows the revenues gained by both the airlines due to spill-in. It is clear that Airline 2 spills in fewer Class 5 and 6 passengers as it moves to DP from EMSRb which results in Airline 1 losing some of the revenues from these passengers. Airline 1 spills fewer Class 4 and 5 passengers into Airline 2 as Airline 2 moves to DP.

We also supplement our findings by looking at the closure rates as shown in

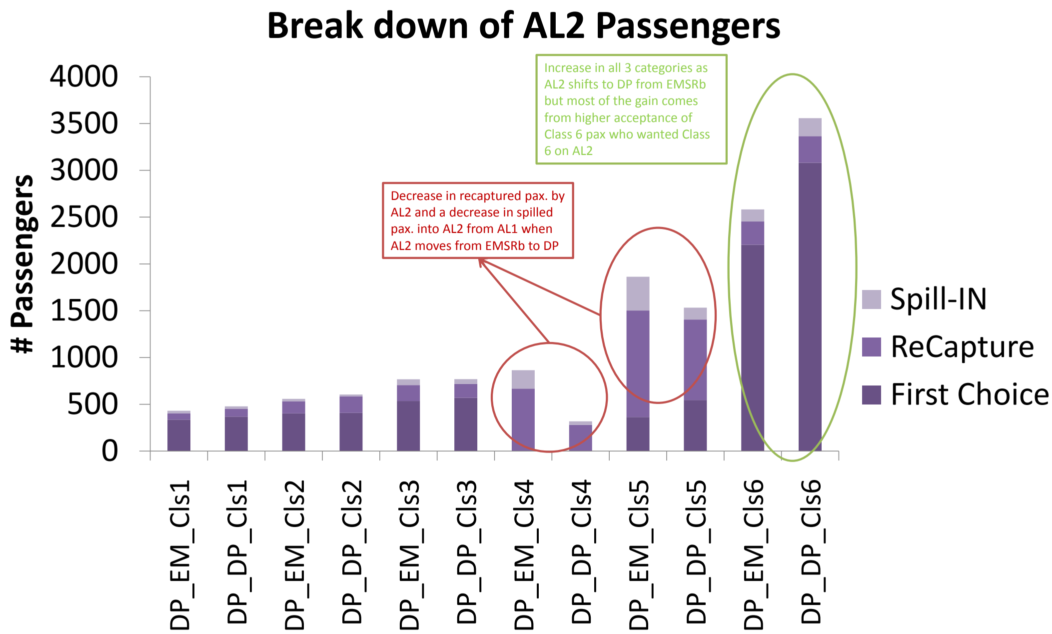

Figure 5-14: AL2 Pax Breakdown

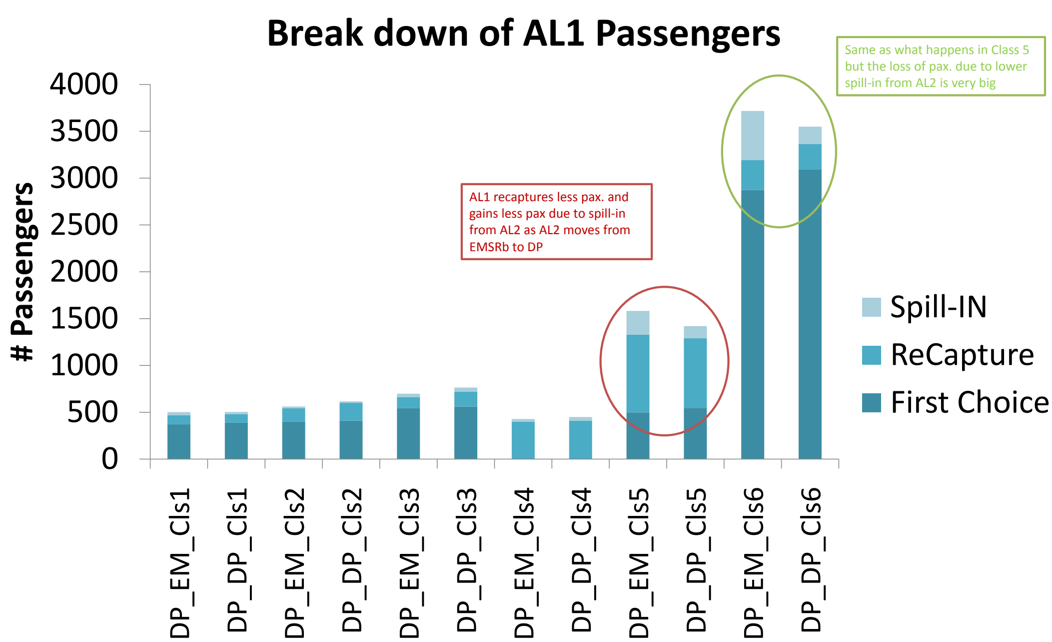

Figure 5-15: AL1 Pax Breakdown

Figure 5-17. Airline 2 is more open in Class 6 throughout and hence accepts more bookings when it uses DP as against EMSRb. Airline 2 is more closed than before in Class 4 and hence it fails to recapture the demand (We also saw this in the Fare

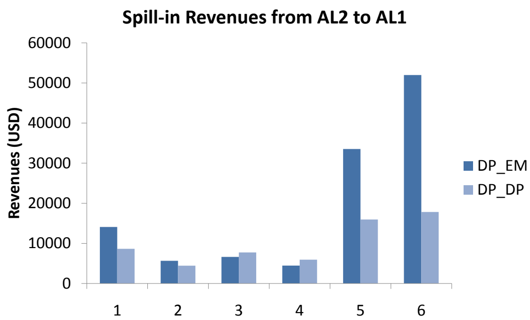

(a) Airline 1

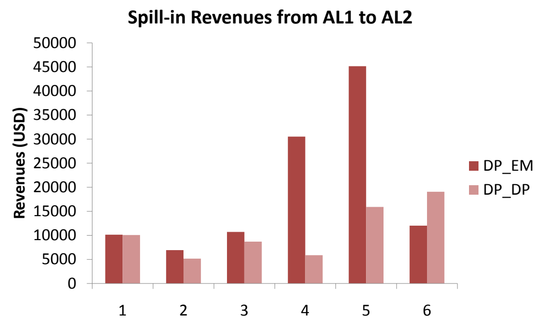

(b) Airline 2

Figure 5-16: Spill-in Revenues in (a) AL 1 and (b) AL 2

Class Mix Graph).

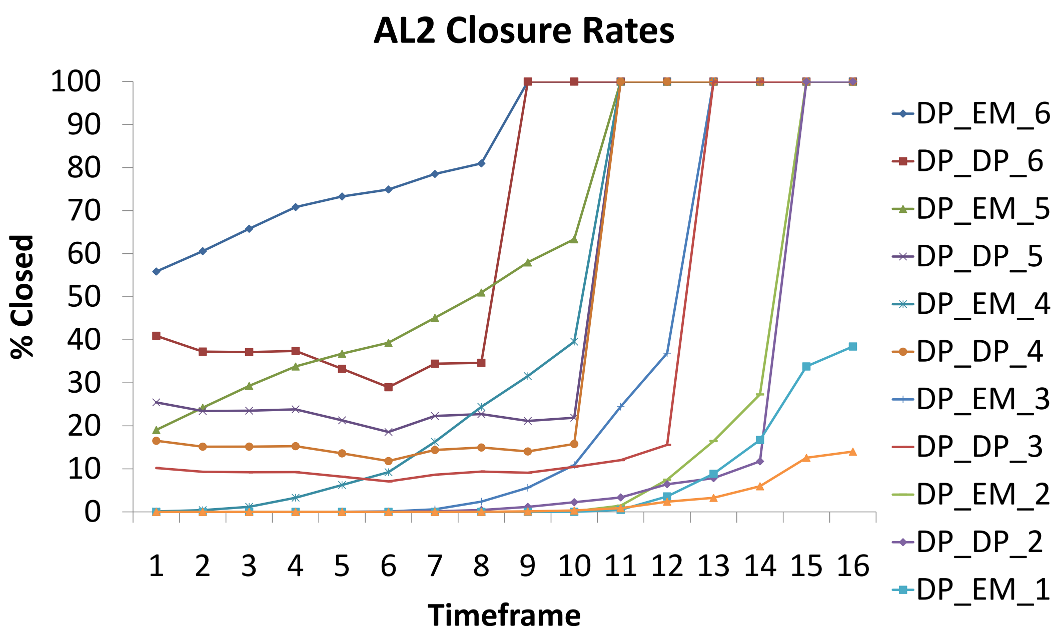

Figure 5-17: AL2 Closure Rates

Table 5.4 reports the No-Go Rates and the Spill-out rates as percentage of total passengers. We observe that both the No-Go rate and the Spill-Out rate decrease as Airline 2 switches from EMSRb to DP. We saw in the closure rates that Airline 2 is less closed in these classes when it uses DP as compared to the case when it uses EMSRb. The above table supplements the argument that more people who wanted Classes 5 and 6 on AL2 were accepted by the airline. Airline 1 lost some passengers in these classes (as seen before from the fare class mix graphs).

	Spill-Out		**No-Go**	
	EMSRb	**DP**	**EMSRb**	**DP**
AL1 Class 5	8.60%	7.20%	5.20%	4.40%
AL2 Class 5	16.10%	6.80%	7.40%	4.10%
AL2 Class 6	12.20%	5.80%	12.70%	12.10%

Table 5.4: Spill-Out and No-Go Passenger Percentage as a Function of AL2 RM System

We discussed at the start of this analysis that we expect fewer lower Class passengers on Airline 1 due to reduced inter-airline spill-in and more Class 6 passengers on Airline 2 due to more availability in Class 6 since Airline 2 uses DAVN-DP which accepts more lower class passengers and fewer higher class passengers resulting in a lower yield but higher load factor. Based on the results and arguments presented till now, we can summarize the following about the passenger movement within the industry between the two scenarios:

- **Airline 1**: Most of the loss in Classes 5 and 6 is due to reduced spill-in from Airline 2 to Airline 1

- **Airline 2**: Airline 2 loses passengers in Classes 4 and 5 due to its failure to to recapture passengers and less spill-in from Airline 1; Class 6 gains are mainly due to more acceptance of Class 6 passengers whose airline preference is Airline 2 since Airline 2 is more open in Class 6 when it uses DP as against EMSRb

Hence we observe that when Airline 2 shifts to DAVN-DP with bidprice control, Airline 1 loses Class 5 and 6 passengers as discussed above. This translates to a loss in revenues for Airline 1. This concludes our discussion on the effect of competition on the performance of DP methods. Now, we turn our attention to the effect of sell-up on DP's performance.

We recall that in the previous set of experimental results with First Choice Only Choice constraint, the network RM methods did not perform better than standard DAVN but with PODS passenger choice, we saw that there is a significant benefit of using DP based methods over standard DAVN with standard forecasting. Apart from the competitive effects discussed above, we discuss the sell-up rates that help DP methods outperform standard DAVN.

The principal idea behind investigating the sell-up rates is that since DP aggressively protects seats for the higher fare classes, the use of DP based optimizers would force more passengers to sell-up to higher classes due to fewer seats available for lower fare classes. This phenomenon would force more passengers to buy fare classes that are closer to their willingness-to-pay and and increase revenues.

		DAVN-DP		
	Std. DAVN	Availability Control	Bidprice Control	UDP
Class 1	-	-	-	-
Class 2	3.10%	5.70%	5.50%	5.50%
Class 3	7.60%	11.40%	10.20%	10.40%
Class 4	0.00%	0.00%	0.00%	0.00%
Class 5	14.40%	13.50%	14.80%	14.70%
Class 6	20.60%	21.30%	21.10%	21.00%

Table 5.5: Sell-up Rates in Network RM

Table 5.5 shows the sell-up rates by fare class for the different optimizers. The figures in the table tell us the percentage of passengers forced to sell-up to higher classes when the passengers' first choice was that fare class and therefore there is obviously no data for Class 1 sell-up rates. We observe from Table 5.5 that all the DP based methods have higher sell-up rates than standard DAVN in all classes. A part of the reason is the aggressive seat protection by DP which makes fewer seats available in the lower classes, thereby, forcing passengers to sell-up to higher fare classes which are more open as compared to standard DAVN.

Table 5.6 shows the revenues due to sell-up by fare class for all the optimizers used in this experiment. As expected, more sell-up to various classes shown by Table 5.5 translates to larger revenues due to sell-up in each fare class when DP methods are used. The table shows us that the additional revenues due to sell-up are much more when DP methods are used as against the case when standard DAVN is used.

		DAVN-DP		
	Std. DAVN	Availability Control	Bidprice Control	UDP
Class 1	9,118	12,587	10,723	10,826
Class 2	15,773	19,632	16,039	16,542
Class 3	4,796	6,507	4,849	5,054
Class 4	17,871	18,154	16,023	17,818
Class 5	11,907	10,450	13,689	12,633
Class 6	0	0	0	0

Table 5.6: Sell-up Revenues in Network RM

Summary

The single flight-leg simulation study done in Chapter 3 did not account for passenger sell-up, inter-airline spill-in and the effects of competition. Those results suggested that though DP is theoretically appealing for the choice of an optimizer in RM systems, the benefit of using DP models over EMSRb is very small unless the realized demand variability is low, or the fare ratios are low, or the capacity is small relative to mean demand, or the passenger arrival pattern can be assumed to be uniform across various fare classes and time frames. However, the results from the experiments with standard forecasting and PODS Passenger Choice suggest that DP based methods outperform standard EMSRb (in leg-based RM) and standard DAVN (in network-RM systems) in most of our simulations. The only exception is the leg based DP with availability control which underperforms relative to standard EMSRb.

We found that the benefits of using DP based methods in RM systems can be attributed to its more aggressive seat protection policy and more passenger sell-up relative to standard EMSRb. As seen in the results from single flight-leg and in Network D6, DP with availability control accepts fewer class 6 bookings because it makes fewer seats available for class 6 in the initial time frames. Although DP opens up more class 6 seats in later time frames, all the experiments show that DP falls short of standard EMSRb in terms of class 6 passenger bookings. We conclude that DP with bidprice control generates more revenues than DP with availability control.

We discussed the effect of aggressive seat protection policy on competitive effects like inter-airline spill-in which aids DP in generating incremental revenues over the base case in both leg-based and network-based RM systems. We also discussed that higher sell-up rates in DP methods compared to standard DAVN helps DP in generating additional revenues. However, the gains of using DP based methods decrease when competitor switches to a more sophisticated RM approach due to competitive feedback effects.

5.2.3 Hybrid Forecasting

In the previous section, we presented various simulation results under different competitive scenarios and simulated the use of various leg-based and network RM techniques for Airline 1. The forecasting method used was standard pickup forecasting method. In this section, we discuss the results obtained by using Hybrid Forecasting as the forecasting technique. We expect that the use of Hybrid Forecasting will lead to a gain in revenues under this fare structure.

Competitor Uses EMSRb

As with standard forecasting, we first present results when the competitor uses standard EMSRb. Figure 5-18 shows the revenues and load factors for leg RM systems

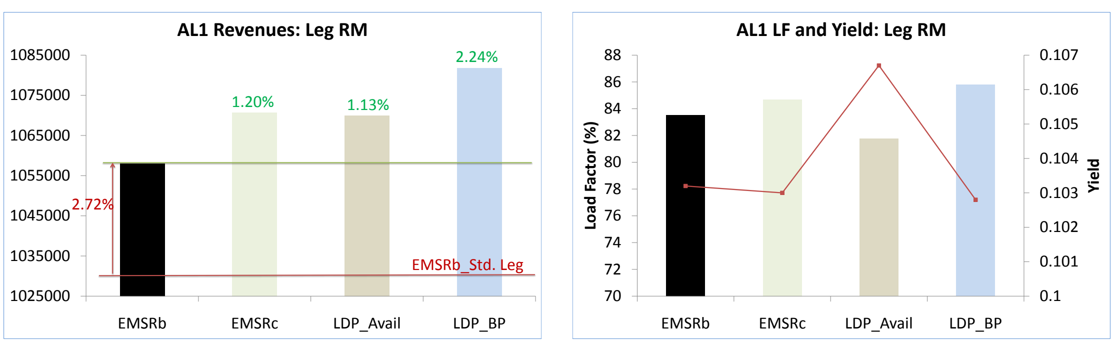

(a) Revenues (b) Load Factors

Figure 5-18: Revenues and Load Factors (a) Revenues and (b) LF and Yield

with Hybrid Forecasting (HF). We observe that a move from standard leg forecasting to hybrid forecasting increases EMSRb revenues by 2.7%. The leg RM techniques generate additional 1.20%–2.25% revenues over EMSRb with hybrid forecasting. As with standard path forecasting, DP with bidprice control outperforms all other optimizers studied in this experiment.

Unlike standard path forecasting where DP with availability control lost revenues relative to standard EMSRb, it outperforms EMSRb with hybrid forecasting. We look at load factors and observe that the difference in load factors generated by both the methods have reduced considerably. The fare class mix shown in Figure 5-19 shows us that the number of Class 6 passengers accepted by DP with availability control are not hugely different from other optimizers when hybrid forecasting is used. We can compare the maximum difference in Class 6 bookings between LDP with availability control and other optimizers. This difference was about 1800 with standard forecasting while it is reduced to about 800 with hybrid forecasting. Figure 5-20 compares the Class 6 closure rates with DP/availability control with the two forecasting techniques. It can be clearly seen that with hybrid forecasting the airline is more open in Class 6 which leads to an increased Class 6 bookings as shown in the fare class mix and hence, an increase in revenues which aids DP/availability control in outperforming EMSRb.

Figure 5-21 shows the revenues and load factors for network RM systems. We observe that hybrid forecasting increases DAVN revenues by about 3.5% over standard DAVN. The increase in revenues obtained by all the network DP methods is small but consistent. As seen before, using network DP methods increases the airline's load factors while decreasing yields.

The fare class mix graph can be used to explain the load factor and yield graph. Figure 5-22 shows the fare class mix for airline 1. As with standard forecasting, the DP methods book much more Class 6 passengers as compared to DAVN whereas DAVN books the most Class 1 passengers. The higher load factor and lower yield for DP methods as seen in Figure 5-21 is a result of DP methods booking more Class 6 passengers and slightly fewer Class 1 passengers which increases the load factors but reduces the yields.

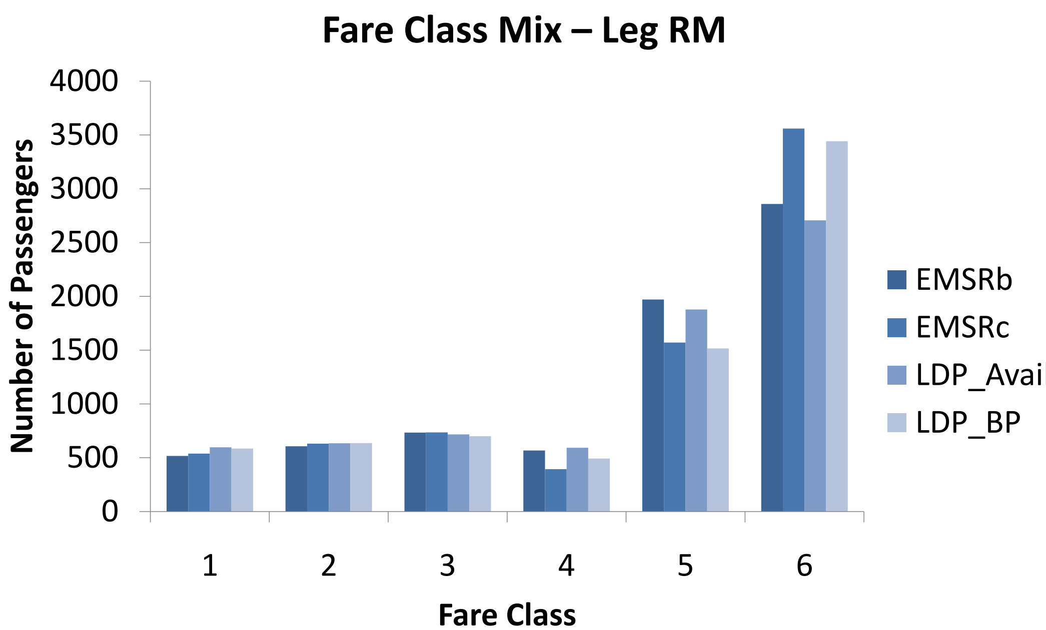

Figure 5-19: Fare Class Mix for Airline 1

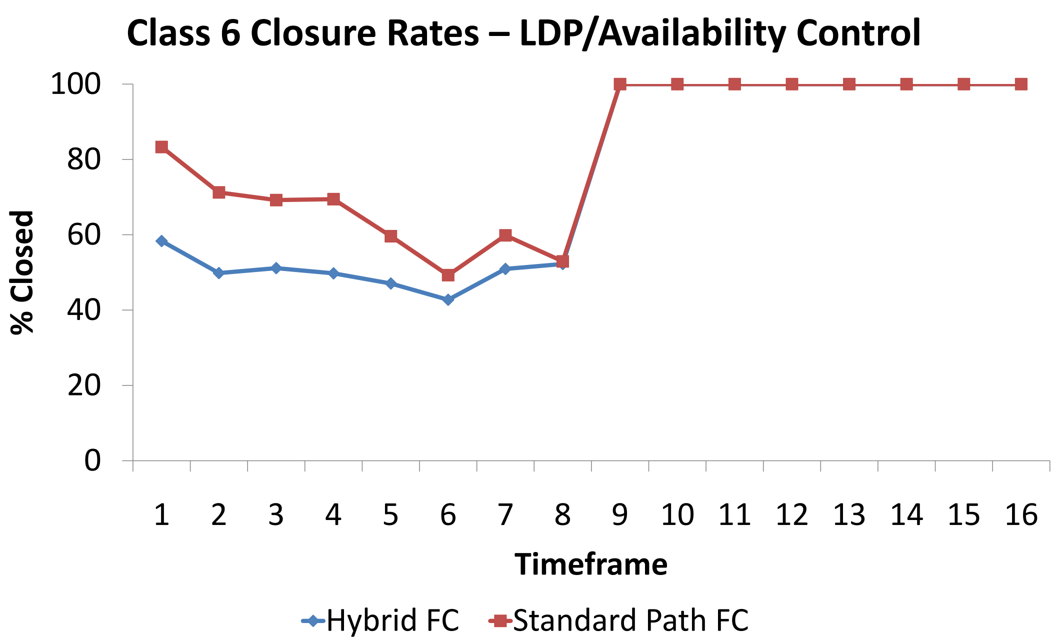

Figure 5-20: Fare Class Mix for Airline 1

Competitor Uses Same RM System as Airline 1

In this set of experiments, competitor (Airline 2) uses the same optimizer that Airline 1 uses. As with standard path forecasting, this set of results will enable us to identify

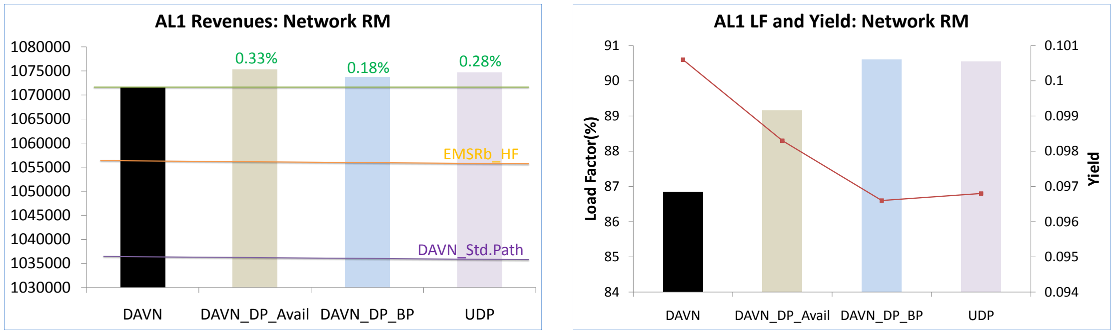

(a) Revenues

(b) Load Factors

Figure 5-21: Revenues and Load Factors (a) Revenues and (b) LF and Yield

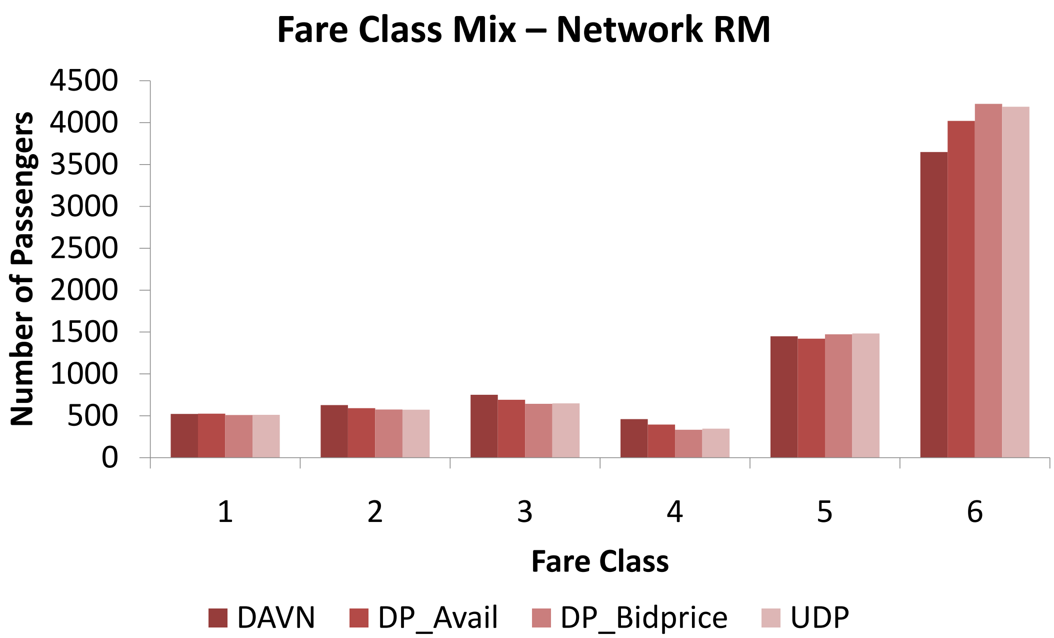

Figure 5-22: Fare Class Mix for Airline 1

the effects of competition on the performance of DP based optimizers.

Figure 5-23 shows the revenues and load factors for leg RM systems. We can see that though DP methods outperform EMSRb, the gains are much smaller than the case when Airline 2 uses EMSRb. As seen before, these results show us that some of the revenue gains of using DP against a competitor using EMSRb is due to DP's aggressive seat protection policy, as was the case with standard forecasting. As airline 2 moves to a more aggressive optimizer, the revenue gains by using DP get

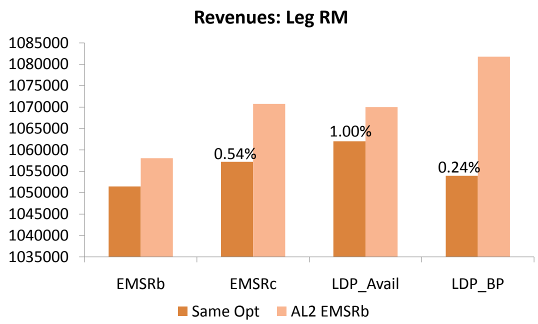

(a) Revenues

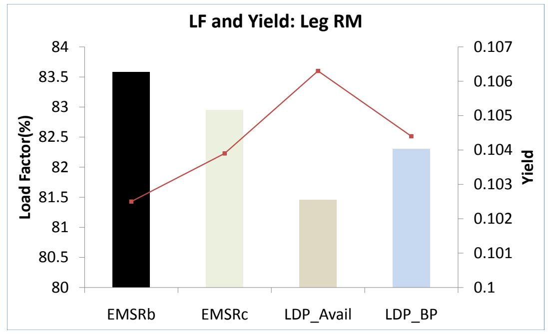

(b) Load Factors

Figure 5-23: Revenues and Load Factors (a) Revenues and (b) LF and Yield

smaller. The load factor and yield graphs are qualitatively the same as with standard forecasting.

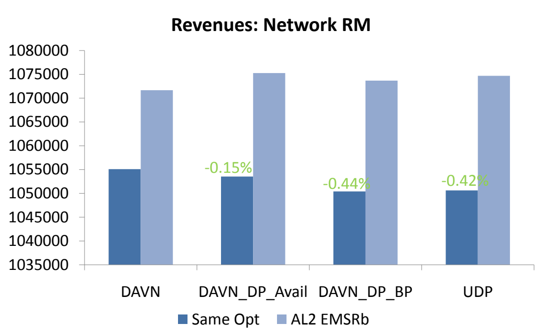

(a) Revenues

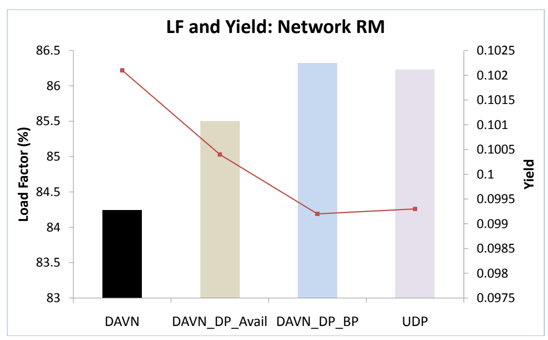

(b) Load Factors

Figure 5-24: Revenues and Load Factors (a) Revenues and (b) LF and Yield

Figure 5-24 shows the revenues and load factors for network RM systems. We observe the same effect of competition — a more aggressive competitor optimizer reduces the benefits of using DP by Airline 1. We also observe for the first time that all the DP methods lose revenues relative to DAVN with hybrid forecasting and symmetric RM systems. The load factors increase and the yields decrease relative to DAVN when Airline 1 uses network DP methods which is consistent with our previous observations.

Summary

The main theme in our results from PODS simulations till now has been that though DP offers theoretically attractive alternative to EMSRb as the optimizer in an airline RM system, the results from single flight leg simulations show that it does not outperform standard EMSRb by a significant margin. The results from PODS, on the other hand, show that in a competitive network simulation DP outperforms standard EMSRb as shown in the previous section with standard path forecasting. We discussed inherent properties of DP like its aggressive seat protection policy which help DP in a competitive setting as DP gathers more passengers due to passenger sell-up, and competitive effects like the inter-airline spill-in.

The results from hybrid forecasting show the same results as with standard path forecasting. DP performs better than the base cases in leg RM and network RM simulations when the competitor uses standard EMSRb. The gains obtained from using DP are higher than those obtained when DP is used with standard path/class forecasting. This makes intuitive sense since hybrid forecasting is a forecasting technique designed specifically to improve the RM system's performance in semi-restricted and unrestricted fare structures.

The use of hybrid forecasting on standard EMSRb alone increases revenues by close to 2.75% and leg-based DP methods increase revenues by another 1.25%-2.25% on top of EMSRb with hybrid forecasting. We obtain similiar results in network RM systems. Hybrid forecasting alone increases revenues by about 3.5% over standard DAVN and hence additional revenues that DP can generate over DAVN with hybrid forecasting are smaller (0.20%–0.35%). as with standard path/class forecasting, we observe that DP methods in network RM accept more Class 6 passengers as compared to DAVN. We observe that the inherent aggressiveness of DP and that of hybrid forecasting increases revenues over standard EMSRb by about 4%.

An interesting point to observe here is that the effect of aggressiveness of RM optimization/forecasting techniques is not additive which can be seen from the network RM simulation results. The incremental revenues of using aggressive techniques like network RM, DP, and hybrid forecasting do not equal to the marginal revenues obtained by using each of these techniques in isolation by themselves. As a matter of fact, we observe this effect in a much more exaggerated manner when Airline 2 moves to a more sophisticated RM system (DAVN-DP) and Airline 1 loses revenues relative to base case by using DAVN-DP and hybrid forecasting. As observed in the previous case with standard path forecasting, the gains of using DP decrease as Airline 2 moves to a more sophisticated RM approach. In leg RM simulations, despite this reduction in revenues for Airline 1, DP methods still outperform EMSRb with hybrid forecasting. However, in network RM simulations, too much aggression due to network RM, DAVN-DP and hybrid forecasting leads to Airline 1 losing revenues when it uses DP relative to DAVN with hybrid forecasting.

5.2.4 Fare Adjustment

In the previous chapter, we presented the concept of fare adjustment. We test the performance of DP methods with Hybrid Forecasting and Fare Adjustment in this section. We expect that the use of Fare Adjustment along with Hybrid Forecasting will lead to a gain in revenues.

Competitor Uses EMSRb

Figure 5-25 shows the revenues for all the optimizers under three different scenarios – no fare adjustment is used, fare adjustment with a scaling factor of 0.50 is used and fare adjustment with a full scaling factor of 1.0 is used. We observe that fare adjustment leads to a gain in revenues for all the optimizers studied in this thesis. DAVN-DP with availability control benefits the most when fare adjustment is applied and gains 1.2% revenues over DAVN with hybrid forecasting. We obvserve that higher the scaling factor for fare adjustment, the higher the gains due to fare adjustment. But we observe that there is one exception to this rule – UDP. We observe that fare adjustment hurts UDP performance and UDP with a full scaling factor of 1.0 underperforms compared to DAVN with hybrid forecasting. We will examine the causes of this behavior of UDP in the next section.

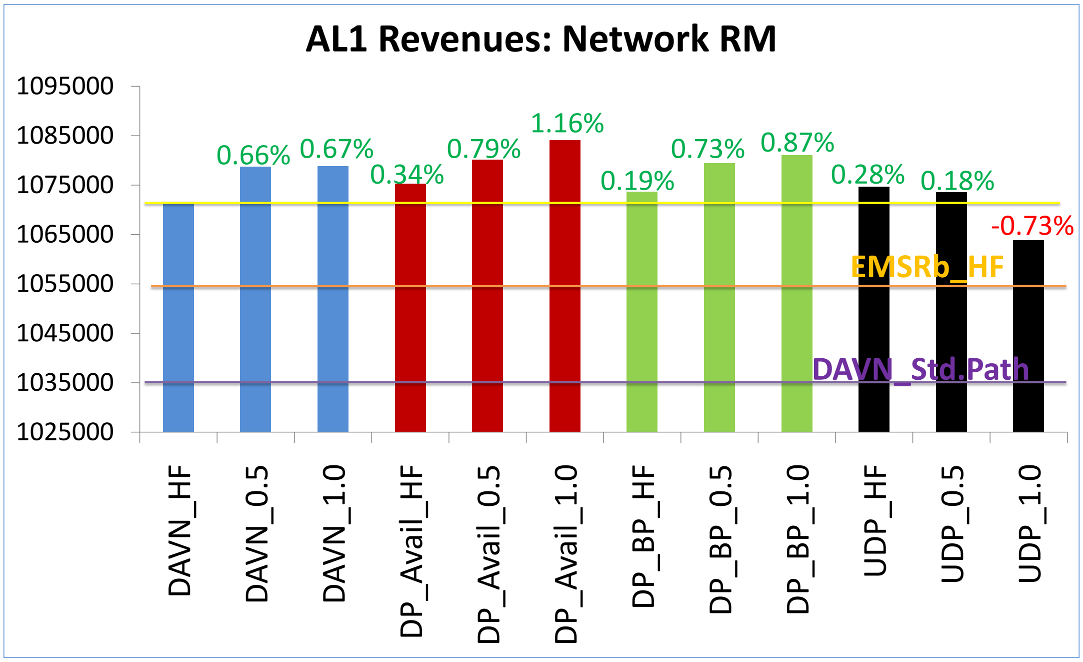

Figure 5-25: Airline 1 Revenues

Figure 5-26 shows the load factors and yields for optimizers with fare adjustment. We observe that the load factors drop as the airline gets aggressive with fare adjustment (as it increases the scaling factor) and the yields increase with an increase in scaling factor.

Competitor Uses Same RM System as Airline 1

Figure 5-27 shows the revenues for all the optimizers as before. We observe that fare adjustment with a full scaling becomes too aggressive in this case and some optimizers lose revenues compared to the baseline of DAVN-HF. We also observe an important change here — network DP methods with a scaling factor of 0.50 outperform the base case of DAVN with hybrid forecasting while DP methods underperform relative to DAVN-HF without fare adjustment as seen in the previous section. DAVN-DP with availability control benefits the most when fare adjustment and the revenue gains increase with an increase in scaling factor. We observed that when Airline 2 uses EMSRb, higher the scaling factor for fare adjustment, higher the gains due to

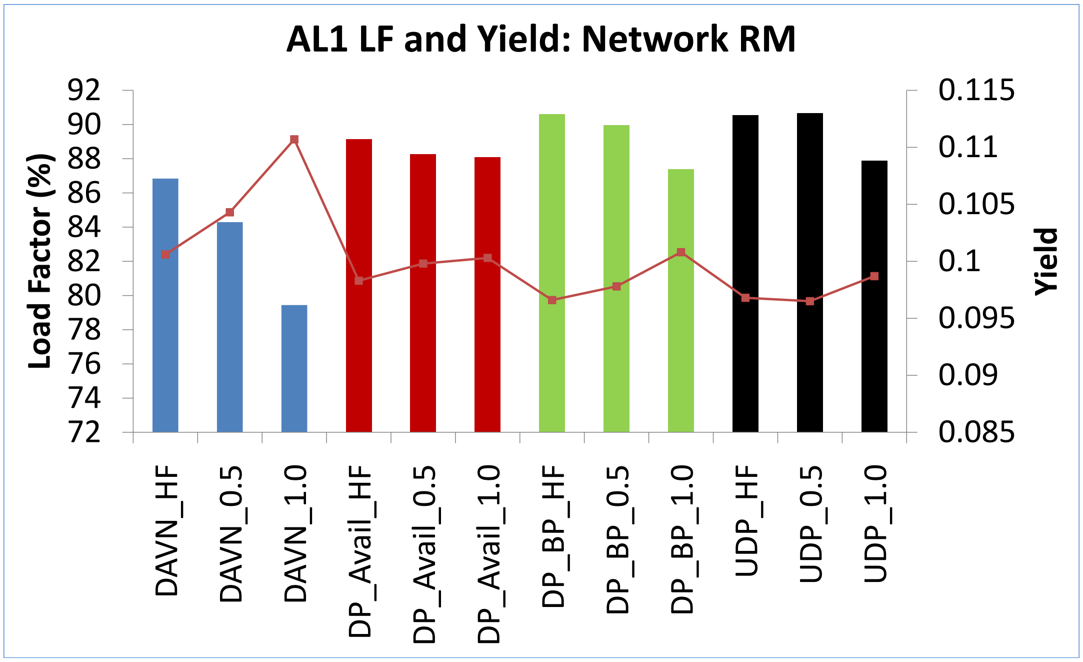

Figure 5-26: Airline 1 Load Factor and Yield

fare adjustment but in this aggressive environment we observe that choosing a right scaling factor dictates the performance of fare adjustment. For most optimizers a scaling factor of 0.50 works better than a scaling factor of 1.0. We see that though fare adjustment with a scaling of 0.50 increases UDP revenues, it still falls short of the baseline.

It is interesting to see that Fare Adjustment with UDP does not work as well as with DAVN-DP. This result was consistent across the experiments that we reported before. Figure 5-29 reports the closure rates of DAVN-DP with fare adjustment and UDP with fare adjustment. We look at this graph to observe the main difference between UDP and DAVN-DP when used with fare adjustment. We can see that UDP gets too aggressive relative to DAVN-DP – UDP is more closed than DAVN-DP except for lower classes in the initial time frames. This leads to a loss in revenues as seen above.

We supplement our argument with the fare class mix graph as shown in Figure 5-29. UDP accepts much more Class 6 passengers since it is more open in the initial Time frames and Class 6 passengers tend to arrive early in the booking process. But

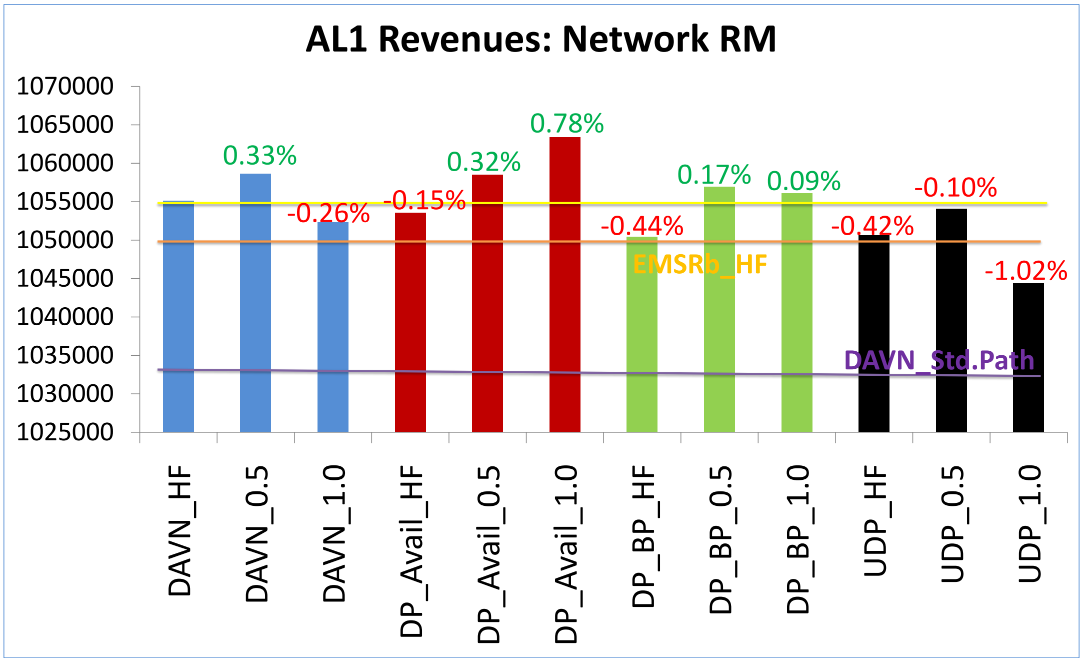

Figure 5-27: Airline 1 Revenues

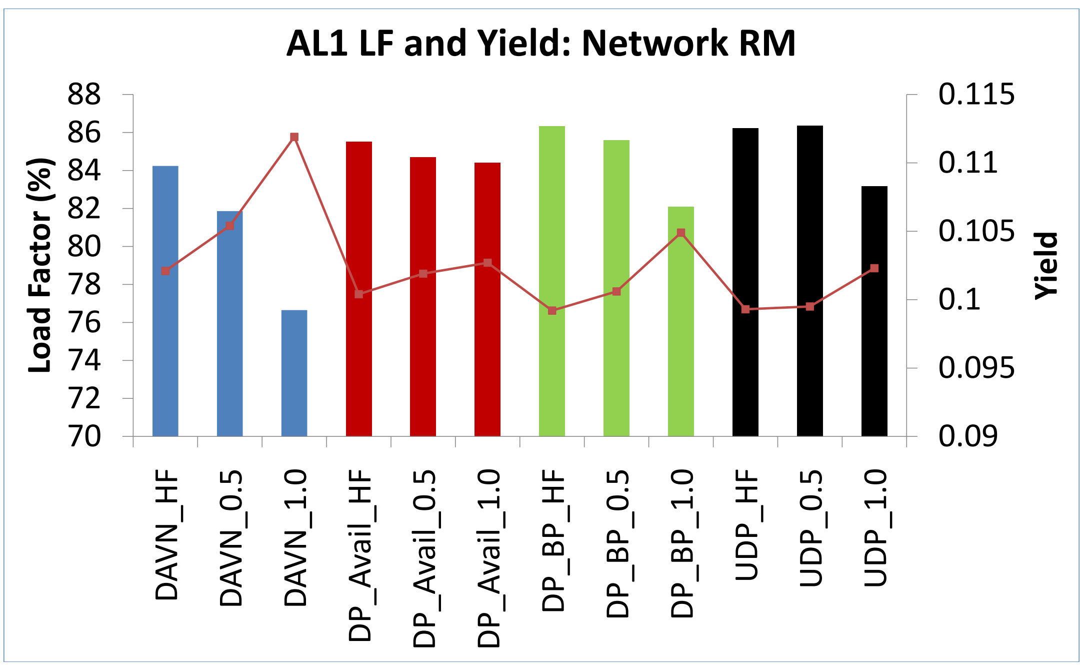

Figure 5-28: Airline 1 Load Factor and Yield

UDP loses a lot of Class 2, 3, 4 and 5 passengers due to higher closure rates as seen above which leads to a loss in revenues while DAVN-DP gains revenues with fare adjustment.

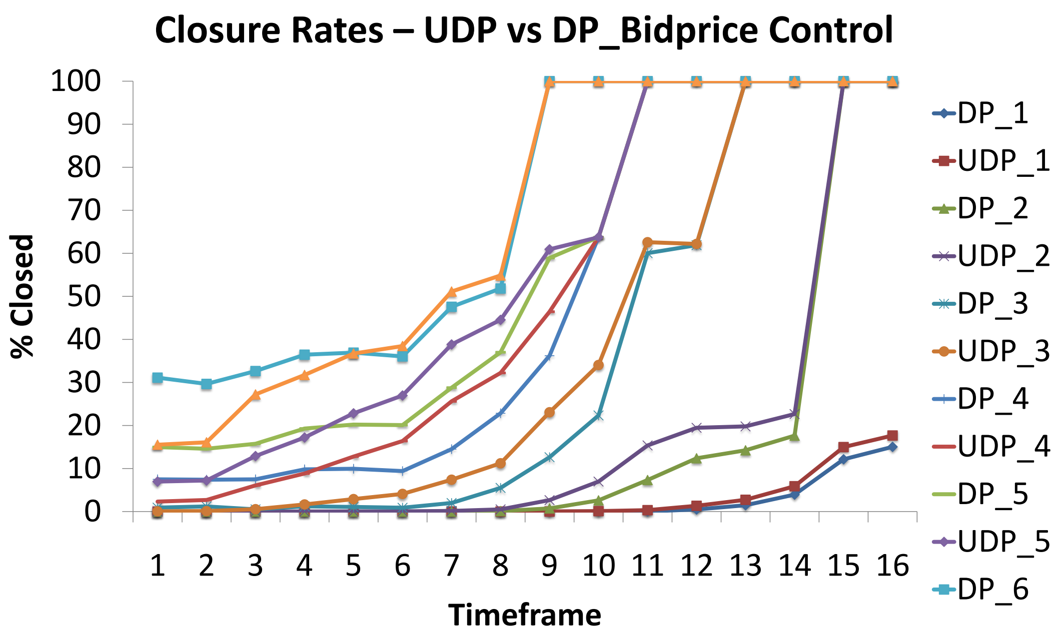

Figure 5-29: Airline 1 Closure Rates

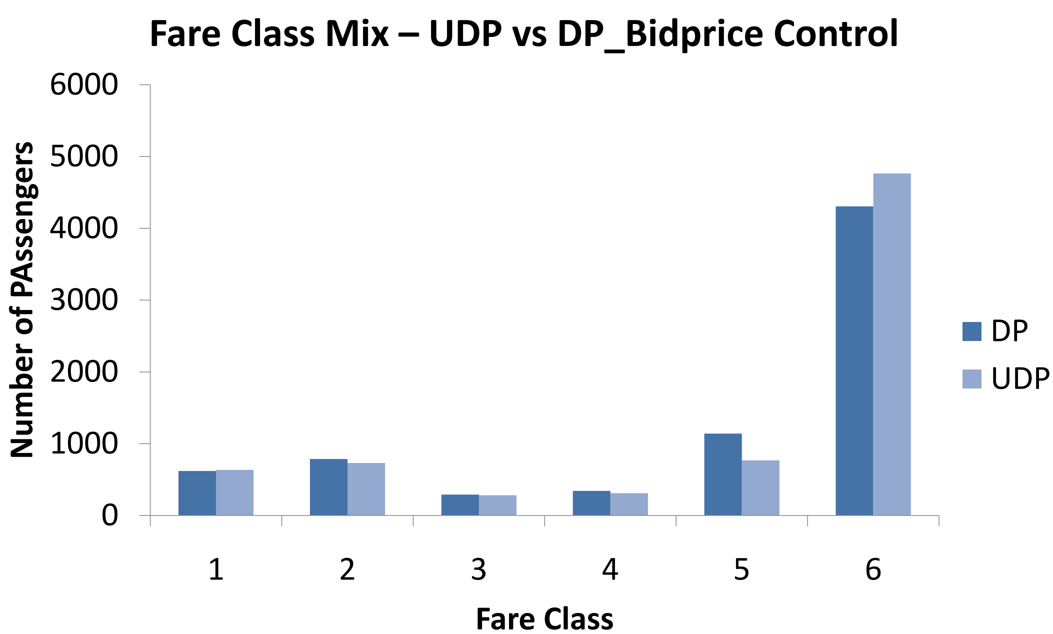

Figure 5-30: Airline 1 Fare Class Mix

Summary

We can summarize our findings with fare adjustment as follows:

- Fare Adjustment increases revenues when AL2 uses EMSRb (except with UDP)

  - As scaling increases, the gain due to FA increases

- The gains due to FA decrease substantially when AL2 uses the same RM method as AL1

  - But FA with right scaling (0.50) increases revenues

  - FA with full scaling performs worse than a scaling of 0.50 in terms of revenues except with DAVN-DP and Avail. Control

- FA decreases load factors and increases yields

- UDP with FA gets too aggressive and closes down fare classes more than DAVN-DP with BP control does with FA

  - Loses many Classes 2-5 passengers and the gain in class 1 and 6 passengers is not sufficient to offset the loss of revenues from classes 2 through 5

5.3 Fully Restricted Fare Structure

We discussed the performance of DP methods under a semi-restricted fare structure in the previous section. In this section, we present some results for a fully restricted fare structure in the network D6. The fare structure along with fare values and restrictions by class are shown in Table 5.7.

Fare Class	Avg. Fare	AP	R1	R2	R3
1	412.85	0	0	0	0
2	293.34	3	0	1	0
3	179.01	7	0	1	1
4	153.03	14	1	1	1
5	127.05	14	1	1	1
6	101.06	21	1	1	1

Table 5.7: Network D6 Fully Restricted Fare Structure

We expect that the absolute revenues obtained by all the optimizers will increase under this fare structure. We also expect that DP will perform better under this fare structure as the demand will decrease for lower classes and will increase for higher classes. This demand shift should help DP because of its aggressive protection policy for higher class demand. We will revisit this argument in details when we discuss the results.

The experimental set-up is same as the one with semi-restricted fare structure. For the sake of completeness, we present the experimental set-up in Table 5.8. Both the airlines use standard pickup forecasting and booking curve detruncation. As in

	Leg RM		**Network RM**	
	Airline 1	Airline 2	Airline 1	Airline 2
Base	EMSRb	EMSRb	DAVN	DAVN
I	EMSRc	EMSRb	DP/Availability Control	DAVN
II	DP/Availability Control	EMSRb	DP/Bidprice Control	DAVN
III	DP/Bidprice Control	EMSRb	UDP	DAVN

Table 5.8: Experimental Set-up

the case with semi-restricted fare structures, we present the results for two scenarios — Firstly, First Choice Only Choice and secondly, with a full PODS Passenger Choice model.

5.3.1 First Choice Only Choice

Figure 5-31(a) shows the revenues, load factors and yields for various leg-based optimizers studied in this thesis. We observe that EMSRc outperforms all other optimizers as was shown in the semi-restricted case. The DP methods also outperform standard EMSRb with bidprice control, yet again, outperforming the availability control when used with DP. We also observe the same trend in load factors and yields as was seen in the semi-restricted case. The important thing to note here is that the revenue gains of DP over standard EMSRb are higher than what we observed in the semi-restricted fare structure. The fare class mix for EMSRb and DP with availability control is shown in Figure 5-32 for the two fare structures. It can be seen that a move from semi-restricted fare structure to a fully restricted fare structure translates to a reduction in lower class passenger bookings and an increase in higher class bookings.

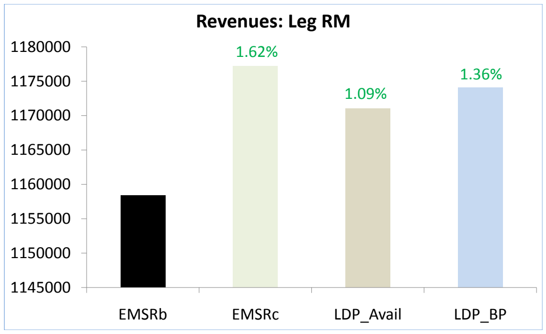

(a) Revenues

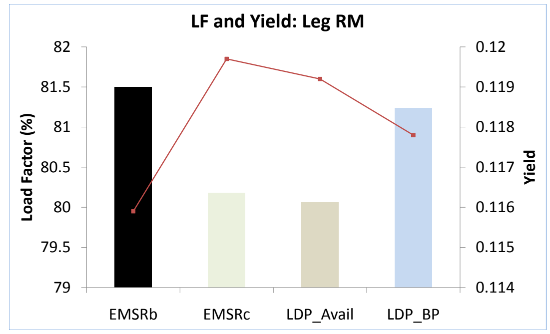

(b) Load Factors

Figure 5-31: Revenues and Load Factors (a) Revenues and (b) LF and Yield

Furthermore, it can be seen that the increase in higher class bookings is greater when Airline 1 uses DP relative to standard EMSRb. We can also observe that the decrease in Class 6 passengers is more for the case when Airline 1 uses DP as compared to standard EMSRb.

In a fully-restricted fare structure, passenger's first choices change due to more restrictions in Classes 4, 5, and 6. One indication of the passenger's choices changing is the abrupt increase in number of bookings for Class 3 as we move to a fully-restricted fare structure. Some of the passengers who would have booked classes 4, 5 and 6 in semi-restricted fare structure find that their utility from purchasing one of these fare classes has dropped substantially due to the restrictions in the fully restricted fare structure. Hence, we see a big increase in class 3 booked passengers since many passengers would find that the total generalized cost[3] of purchasing class 3 is lower than that of classes 4, 5, and 6. We saw in semi-restricted fare structure results that DP with availability control books more higher class passengers and fewer lower class passengers. The revenue gains from greater higher class bookings more than compensates for the losses due to fewer lower class bookings and hence, DP outperforms standard EMSRb. As we have discussed before, DP has an inherent structural property that it aggressively protects seats for higher fare classes. Hence, it is intuitive that DP will perform better when there is high demand for higher fare

[3]Generalized cost is the sum of the price of that fare class and the disutility costs of the restrictions associated with that fare class

classes and lower demand for lower classes which is the case in fully-restricted fare structure as discussed above.

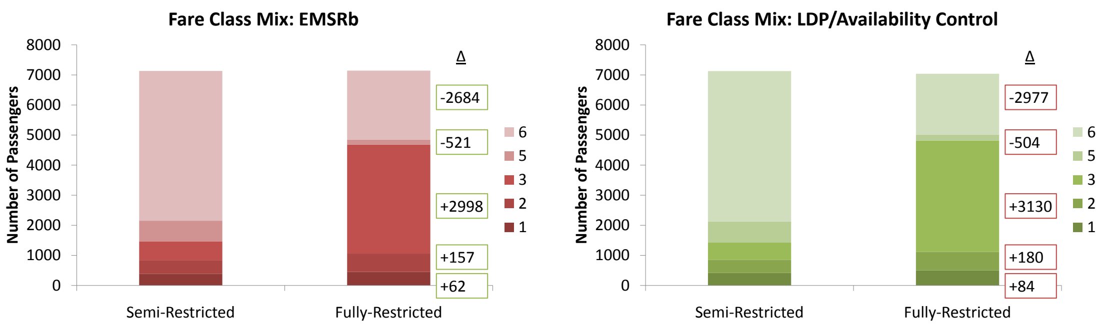

(a) EMSRb

(b) DP/Availability Control

Figure 5-32: Fare Class Mix for (a) EMSRb and (b) DP/Availability Control

The performance of network RM methods is shown in Figure 5-33. We observe that the gain of moving to standard DAVN from EMSRb is 1.22% which is much larger as compared to the semi-restricted results. The three network DP methods outperform standard DAVN by a small margin (0.06%–0.17%) although these gains are bigger than the gains in semi-restricted fare structure. We have already discussed the effects of adding restrictions on the performance of DP and EMSRb while discussing the results in leg RM and those points hold good in network RM systems also. The load factor increases due to the use of network DP methods and the yield decreases which is similar to our previous observations in all experiments.

Summary

We observe the same qualitative results as seen and discussed before. Hence, we can summarize our findings from this small experiment as follows:

- Leg RM
  - EMSRc once again does better than any leg RM method
  - DP performs better than EMSRb (gain of 1.1%-1.4%)
    - Bidprice control works better than Availability control

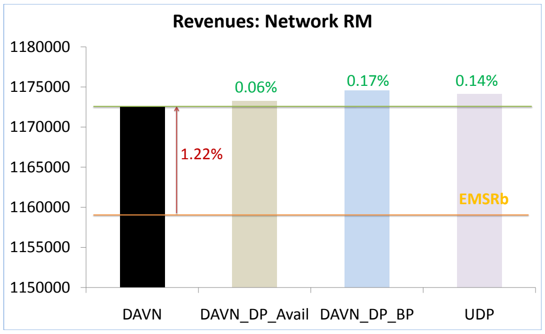

(a) Revenues

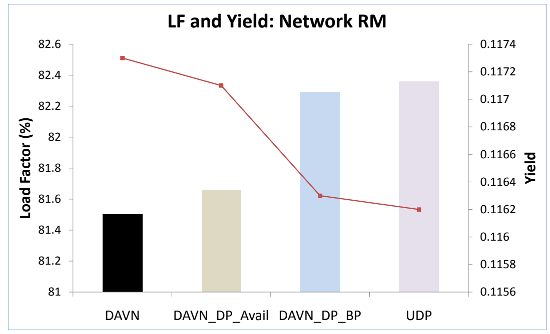

(b) Load Factors

Figure 5-33: Revenues and Load Factors (a) Revenues and (b) LF and Yield

- Network RM
  - The move from leg-based RM to network RM (EMSRb to DAVN-EMSRb) increases revenue by a bigger margin than before
    - Hence The gain of using DP over and above DAVN is much less than previous results
  - Bidprice control works better than availability control
- The benefits obtained by DP exclude the effects of sell-up and competition since it is First Choice Only Choice

5.3.2 PODS Passenger Choice

In the previous section, we presented and summarized the results from FCOC experiments under fully restricted fare structures. We now move on to more elaborate result discussions when passenger choice model is used.

Competitor Uses Standard EMSRb

The competitor (Airline 2) uses EMSRb with standard leg/class forecasting in all the experiments and Airline 1 is the airline of interest. Figure 5-34 shows the results for leg RM methods. As seen before under semi-restricted fare structure, incorporating the passenger choice model increases the revenues of all optimizers as well as the revenue gains of all the optimizers over EMSRb. We also observe, as with semi-restricted fare structure, that DP with bidprice control outperforms all other optimizers at the leg level. The load factors and yields for the various optimizers follow similiar trend as in semi-restricted network.

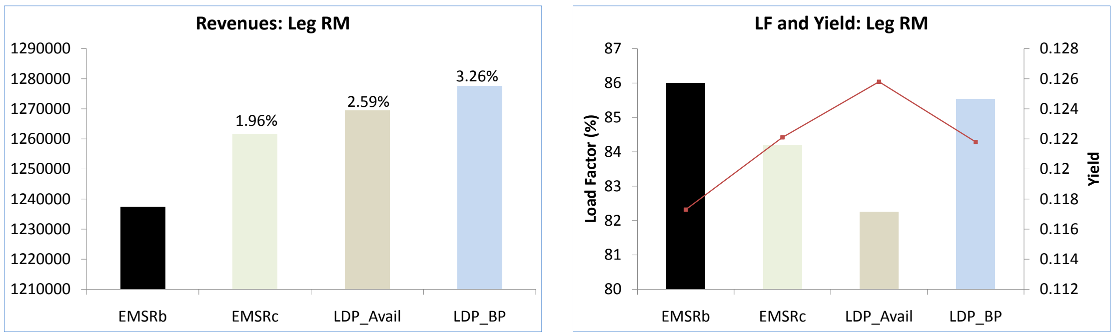

(a) Revenues

(b) Load Factors

Figure 5-34: Revenues and Load Factors (a) Revenues and (b) LF and Yield

As seen with FCOC in restricted fare structure, the revenue gains are much bigger than the semi-restricted network and DP with bidprice control outperforms EMSRb by almost 3.3%. We have already discussed that the additional strong restrictions present in the fully-restricted network increase class 3 bookings by a big margin. In general, the higher classes bookings increase and lower classes bookings decrease due to an increase in generalized cost of the three lowest classes to which restrictions were added. The fare class mix shown in FCOC accurately depicts the passenger's first choice and we showed that fare class mix supports the above argument. With a full PODS passenger choice model, booked passengers include passengers who were forced to sell-up as well as passengers due to inter-airline spill-in. Hence, we look at the average number of passenger demand generated over 5 trials of 400 samples (600 samples with 200 samples burnt) by fare class in the two fare structures to study the effect of adding restrictions on passengers' first choice.

Table 5.9 confirms our intuition that passengers' first choice changed due to addition of restrictions such that demand in higher classes increased and in lower classes

Class	Semi-Restricted	Restricted	Difference	% Difference
1	481	555	75	15.58%
2	486	658	172	35.30%
3	704	3,827	3,123	443.59%
5	801	244	-557	-69.57%
6	5,704	2,892	-2,812	-49.29%

Table 5.9: Average Demand by Fare Class

decreased. These results are in-line with our analysis of the FCOC results and hence, like in FCOC case, the aggressiveness of DP helps it perform better than semi-restricted case (relative to standard EMSRb) because of higher demand in higher classes and lower demand in lower classes as discussed before.

Figure 5-35 shows the revenues and load factors for network RM methods. Recall that the network RM methods did not show any significant benefits with FCOC experiments described before. We observe from Figure 5-35 that the move from EMSRb to standard DAVN increases revenues by close to 1.50%. The network DP methods outperform standard DAVN by 0.20%- almost 1%. The trend seen in the performance of network RM methods is very similiar to the one seen before — Under FCOC, network RM methods do not perform much better than DAVN but with PODS passenger choice, all the network DP methods outperform DAVN. As discussed before, the load factor and yield graph is similar to the one from semi-restricted case in which the load factors increases for the bidprice methods and the yield decreases when used with network DP methods.

Competitor Uses same RM System as Airline 1

As before, to eliminate the effect of aggressive seat protection limits and competition, we present the results with symmetric RM systems in which both airlines use the same optimizers.

Figure 5-36 shows the performance of leg RM methods. As seen before, though DP methods outperform EMSRb, the gains in revenues decrease in most cases as the competitor moves to a more sophisticated RM system. DP with bidprice control, as

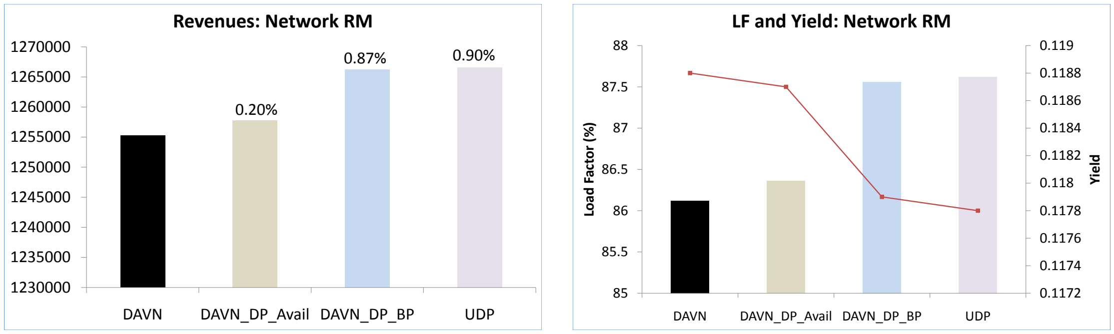

(a) Revenues

(b) Load Factors

Figure 5-35: Revenues and Load Factors (a) Revenues and (b) LF and Yield

before, outperforms all other RM systems and it is the only case when the revenues increase as the competitor moves to a more sophisticated RM system.

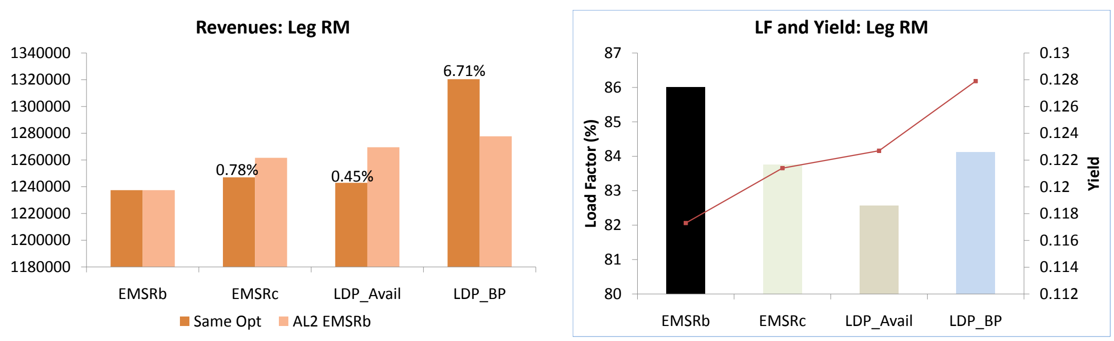

(a) Revenues

(b) Load Factors

Figure 5-36: Revenues and Load Factors (a) Revenues and (b) LF and Yield

Figure 5-37 shows the performance of network RM methods. The results are very consistent with our previous results. As competitor moves from EMSRb to a more aggressive RM system, the revenue benefits of network DP methods reduce. DP still outperforms DAVN but by a smaller margin which shows that the aggressive seat protection policy in DP (or its aggressiveness) helps DP against a competitor using EMSRb. The load factor and yield graphs are also consistent with all the previous results.

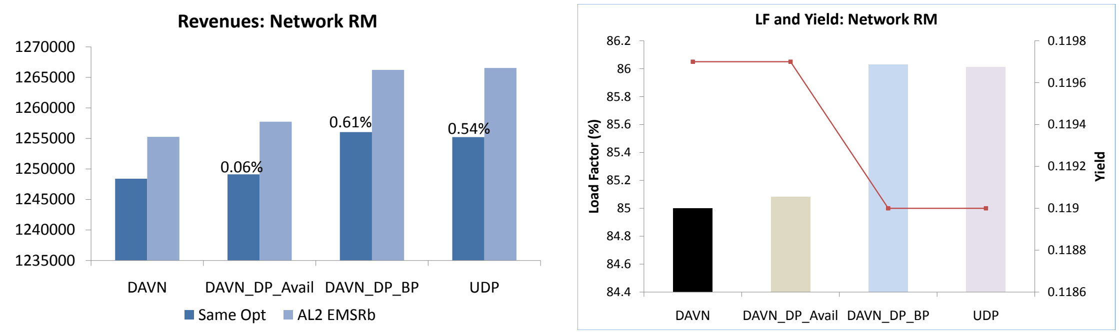

(a) Revenues

(b) Load Factors

Figure 5-37: Revenues and Load Factors (a) Revenues and (b) LF and Yield

Summary

We presented the results in fully restricted fare structure with standard forecasting and found out that the results are consistent with the ones we discussed with a semi-restricted fare structure. Most of the observations made earlier with the semi-restricted fare structure hold under fully restricted fare structure. We observed that the reasons for DP performing better than standard EMSRb is same as before — DP's inherent quality of aggressive seat protection for higher classes, inter-airline spill-in and passenger sell-up. We also saw that the benefits of using DP decrease as the competitor switches to a more sophisticated RM technique which is in line with our previous observations.

We also discussed that the absolute revenues for all the optimizers increase under the fully restricted fare structure because the added restrictions in the lower classes increase the total generalized cost of the classes and hence, lowering the demand for those classes. We observed that passengers' first choice changed from lower classes with the newly added restrictions to the class immediately above them – Class 3. Though this phenomenon was present in both DP and EMSRb, the shift in passenger demand from lower classes to higher classes and the inherent aggressiveness present in DP helped it outperform EMSRb by a bigger margin in the restricted fare structure as against the semi-restricted fare structure.

5.4 Summary

We have seen that DP does not consistently outperform EMSRb in the single flight leg simulation as presented in Chapter 3 despite being the theoretical model of choice for an optimizer in an airline RM system. The main theme of the results and discussions presented in this chapter has been that though DP does not do better than EMSRb in a simplified simulator with no competition and network effects, DP does outperform EMSRb fairly consistently in our PODS simulations. We tested the performance of DP relative to EMSRb in a competitive network setting with the optimizer in the RM system being the only source of difference between the two airlines' operations over this network.

In the semi-restricted fare structure, we observed that DP outperforms EMSRb in almost all leg RM as well as network RM simulations and these results are consistent across FCOC experiments and PODS passenger choice experiments. We also observed that DP with availability control at a leg level protects aggressively for class 1 passengers due to its Poisson assumption (which means it assumes demand variance equal to the mean demand) and hence, makes fewer seats available for class 6 passengers in initial time frames. These results are consistent with our observations in the single flight leg simulations performed in Chapter 3. DP with bidprice control generated greater revenues than DP with availability control throughout the simulations with standard path/class forecasting in this fare structure.

We also observed that with the full PODS passenger choice model, DP generated more revenue gains than with First Choice Only Choice. The discussions demonstrated that the revenue gains are due to competitive feedback effects like inter-airline spill-in and passenger sell-up. We observed that due to its inherent aggressive seat protection for higher fare classes, DP forces more passengers to sell-up to higher fare classes which translates to an increase in revenues. However, when the competitor switches to a more sophisticated RM system (In our experiments, we used s symmetric RM system), we observe that the revenues obtained by DP decrease in most experiments indicating that most of the gains of DP obtained when Airline 2 uses standard EMSRb is due to its inherent aggressive seat protection policy for higher classes.

The results from the experiments done with Hybrid Forecasting again show that DP outperforms EMSRb (or DAVN) in both leg RM and network RM system simulations and the gains from using DP decrease as Airline 2 moves from standard EMSRb to a more sophisticated RM system. However, in network RM simulations, network RM with hybrid forecasting and DP optimizer gets too aggressive and it hurts the performance of DP relative to DAVN. It is also interesting to note that unlike standard path/class forecasting where bidprice control works better than availability control, DP with availability control outperforms DP with bidprice control in most experiments.

All these results indicate that though aggressiveness helps in an RM system, it is not an additive effect and too many aggressive components (network RM, Hybrid Forecasting, DP) do not always make a better RM system. We also tested the DP optimizer with Fare Adjustment. The use of Fare Adjustment increases the revenue gains of DP compared to DAVN. However, UDP loses revenues when used with Fare Adjustment and we discussed that the main reason behind this is that UDP with fare adjustment gets too aggressive and closes down some fare classes which DAVN-DP with bidprice control does not. Our insight from Hybrid Forecasting results that too much aggressiveness hurts RM system's performance is consistent with the Fare Adjustment experiments with UDP. We also observe that the move to a more sophisticated RM system hurts the performance of DP in presence of Fare Adjustment, as was the case in every experiment we discussed.

In the fully restricted fare structure, we observed that DP outperforms EMSRb in all experiments independent of the competitor's RM system and the passenger choice model used (FCOC or PODS passenger choice) although competitor's shift to a more sophisticated RM system reduces the gains of DP over EMSRb (or DAVN). We observed that the absolute revenues of all the optimizers increased under fully restricted fare structure. This is intuitive since the addition of restrictions makes the fare structure fully differentiated and the total generalized cost of lower fare classes become higher than in semi-restricted fare structure. This increases the demand for Class 3 (the lowest fare class which did not get any restrictions added) by a big amount and increases the demand for other higher fare classes by a smaller amount while decreasing the demand for lower fare classes at the same time.

The reasons for DP outperforming EMSRb (or DAVN) are similar to those discussed above for the semi-restricted fare structure. However, an increase in gains of using DP over EMSRb (or DAVN) in a restricted fare structure can be attributed to the demand shift from lower classes to higher classes. DP, by virtue of its aggressive seat protection policy for higher classes, benefits from this structural change more than EMSRb and hence outperforms EMSRb by a bigger margin than in semi-restricted fare structure. All other insights gained from semi-restricted fare structure results hold good for fully restricted fare structure.

We conclude this chapter by summarizing the major insights gained from the PODS simulations performed in this chapter. We observed that DP protects aggressively for Class 1 seats and makes fewer seats available for lower class passengers. However, DP makes more seats available for Class 6 passengers in later time frames which is consistent with our results from Chapter 3. We note that although DP did not perform much better than EMSRb in the single flight leg simulations done in Chapter 3, our simulation results from PODS show that DP performs better than EMSRb (or DAVN) due to its inherent aggressiveness in seat protection for higher classes which helps DP gain more revenues due to competitive feedback effects like inter-airline passenger spill-in, and passenger sell-up.

Bidprice control generates more revenues than availability control with standard path/class forecasting. However, with Hybrid Forecasting, DP gets too aggressive with bidprice control and DP with availability control performs better. The move to a sophisticated RM system by the competitor reduces the gains obtained by DP indicating that most of the revenue gains obtained by using DP are due to its aggressiveness. Another important insight was that too much aggressiveness in individual components of an RM system might hurt the performance as the effect of aggressiveness is not additive. We also observed that DP has a greater potential for revenue gains relative to EMSRb in a fully differentiated fare structure where there is more higher class demand and less lower class demand which helps DP due to its inherent aggressiveness.

Chapter 6

Conclusion

We have seen the theoretical formulations as well as implementation details of various DP methods in previous chapters. Specifically, we discussed variants of basic DP methods implemented in our simplified single flight leg simulator as well as in Passenger Origin Destination Simulator (PODS). The results from the single flight leg example were shown in Chapter 3 using our simplified simulator and simulation results in a competitive network environment obtained from PODS were shown in Chapter 5. We begin this chapter reviewing the thesis objectives as mentioned in the first chapter and then summarize the main findings of this thesis. We also present some possible future research directions at the end of the chapter.

6.1 Summary of Thesis Objectives

The objective of this thesis was to study the performance of Dynamic Programming (DP) based methods in airline Revenue Management (RM) systems. We discussed that most traditional RM methods assume independent demand for various fare classes along with a predetermined and sequential arrival of demand for fare class bookings, typically lowest classes booking earlier than higher classes in the booking process. The booking limits set by traditional RM methods are based on total remaining demand with no distinction among demand levels in future time frames. These booking limits remain static until the next re-optimization point in the booking horizon. Dynamic Programming (DP) based approaches, although assuming independent demands for fare classes, do not make any a priori assumption on the arrival pattern of various fare classes and use the forecasts by time frame and by fare class to set booking limits. The booking limits set by DP models are a function of bookings in hand and remaining capacity. Hence, even within a time frame, the booking limits may change according to the realized demand within a time frame. The booking limit matrix, which is a function of the number of time frames and bookings in hand (or remaining capacity) in each of them, helps the DP models adjust dynamically to changing passenger booking patterns even before the next re-optimization point. We argued that because DP models relax two out of the three limiting assumptions made in traditional RM models, they offer a theoretically attractive alternative to traditional RM methods.

We also discussed that another motivation behind the work in this thesis is that some studies have shown mixed results with DP applications to airline revenue management. Given the attractiveness of DP models as described above, the observation that DP models do not consistently outperform traditional RM methods is not entirely obvious. Hence, one of the goals of this thesis was to study the effect of various assumptions made in the DP models and their impacts on the performance of DP models. Another goal of the thesis was to study the performance of DP models in RM systems and understand the inherent properties of the DP methods that can help DP models outperform standard RM techniques in competitive simulations.

The underlying principle of all dynamic programming models is that the booking horizon is divided into discrete time periods, each of which are small enough so that the probability of more than one booking request is negligible. The problem is modeled as a finite horizon, discrete-time Markov decision process (MDP) with total number of bookings-in-hand as the state variable. The most common method to solve such a formulation is through backward recursion as was discussed during the literature review. The output generated by most algorithms is a bidprice matrix (or vector) as a function of remaining capacity and bookings-in-hand. The bidprices can be translated into optimal booking limits depending on the type of control (availability control or bidprice control) being used by the airline.

We developed our own simplified simulator in Chapter 3 to study the effects of assumed Poisson variance on the performance of DP models. The main purpose of this single flight leg study was to have an idealized simulator which meets the assumptions of DP to provide us with an idealized, and unbiased test-bed. After studying the simulation results obtained from our simplified simulator, we studied the performance of DP methods in PODS in Chapter 5. The main purpose of our simulations in PODS was to study the effectiveness of DP methods under a competitive network environment within a typical airline RM system. We discussed the simulated results of implementing DP methods in two different network environments, a semi-restricted network and a fully restricted network. As done before, we emphasize that the main aim of this study is not to obtain the sensitivity of the revenue gains with respect to the choice of sell-up estimators, or prove what combination of forecaster, sell-up estimator generates highest revenues with DP models described in this thesis. Rather, we try to understand the reasons behind the observed revenue gains and the impacts of assumptions in DP models on these revenue gains from a broader perspective.

6.2 Summary of Results

In this thesis, we developed our own simulator to compare the performance of a basic DP model relative to EMSRb in a single flight leg. The main goal of this simulator was to test the effectiveness of DP based methods under different scenarios in a simplified simulation environment that excludes the competitive feedback effects of inter-airline spill, passenger sell-up within the airline as well as that of forecasting accuracy. We also developed the capability of simulating different demand variability to test the performance of DP based methods under different demand variances. The simulator matches the assumptions of DP models for the simulated *z-factor* of 1. The simulator helped us perform sensitivity analyses on different assumptions of demand factors (sensitivity to mean demand relative to capacity) as well as fare ratios (sensitivity to assumed fare ratios of the highest to lowest fare classes) for two different types of passenger arrival patterns (sensitivity to assumed demand arrival pattern).

The results showed that DP did not outperform EMSRb consistently, despite DP being the theoretical optimizer of choice. The difference between the revenues generated by both the optimizers was small but DP performed better than EMSRb when the simulator matched DP's assumptions. We observed that as the variance of realized demand increased, the relative gains of DP based optimizer decreased relative to EMSRb and observed that EMSRb generated higher revenues than DP for higher levels of demand variance.

We argued, by way of sensitivity analyses, that the benefits of DP depend on various factors like realized demand variability, fare ratios, capacity relative to mean demand, and the passenger arrival process. DP performed well when the demand factors were high or when the fare ratios were low. The performance of DP was even better when the simulated passenger arrival pattern was assumed to be uniform — probability of a booking request from any class in any time frame was equal.

A common theme across all the results was that DP was aggressive in protecting seats for higher classes and made fewer seats available for the lower classes in the initial time frames, as compared to EMSRb. We argued that as DP assumes a Poisson variance, DP is over-certain for higher class booking requests and hence, over protects for Class 1. However, in later time frames, when the assumed Class 1 booking requests do not arrive, DP makes more seats available for lower classes but the over protection hurts the performance of DP.

We found that for realistic *z-factors* (those observed in actual airline RM systems), there is no significant difference between the performance of DP and EMSRb in the base case of the simulation experiment in a single flight leg. An increase in capacity relative to demand (resulting in lower load factors) or an increase in fare ratios adversely affects DP's performance while a decrease in fare ratios leads DP to outperform EMSRb. We conclude by commenting that high demand variance adversely affects the performance of DP based methods and the benefits of having a dynamic booking limit policy (as in DP) are negated by the fact that the assumed demand variance in DP models is much less than the variance realized in most airline

RM forecasting systems.

We then performed simulation experiments on a more realistic simulator, Passenger Origin Destination Simulator (PODS), in a competitive network environment as a natural extension of our experiments on a single flight leg in a simplified simulator. The simplifications made in the simplified simulator were relaxed in PODS simulations. The PODS simulations were performed in a competitive environment with two airlines operating similar schedules, and hence, the RM system used by the airlines was the primary distinguishing factor. This enabled us to study the network effects and competitive feedback effects.

The simulations in PODS were divided into two different categories — First Choice Only Choice (FCOC) experiments and PODS Passenger Choice experiments. In FCOC experiments, a passenger travels only if the first choice requested by the passenger is available and hence, the passenger is counted as a spilled passenger (or a No-Go) if the passenger's first choice is not available. A consequence of such an experiment is that these experiments exclude the impacts of passenger sell-up, recapture and spill-in within an airline or across airlines. The experiments with full PODS passenger choice use an elaborate form of passenger choice model as described in Section 4.2.1. We believe that these experiments capture the impact of competition and passenger sell-up which were absent in FCOC experiments.

The DP based optimization approaches were divided into two categories: Leg based DP models and Network based DP models. The former category includes the standard DP model on a single flight leg while the latter category, which incorporate the network displacement effects in DP models, was introduced in this thesis. The network based DP approaches include DAVN-DP and Unbucketed DP (UDP) which were discussed in Sections 4.5.4 and 4.5.5 respectively. All the DP methods discussed in this thesis were evaluated using two types of controls – Availability control and Bidprice control.

In contrast to the results obtained in the single flight leg simulation experiment, we observed that DP performed better than EMSRb in most of the experiments in PODS. We observed that DP outperformed EMSRb in almost all leg RM as well as network RM simulations regardless of whether First Choice Only Choice is used or PODS passenger choice is used. We also observed that DP with availability control at a leg level protects aggressively for class 1 passengers and makes fewer seats available for class 6 passengers in initial time frames, which is consistent with the results we obtained from the single flight leg simulations. The bidprice control methods generated more revenues than the availability control in most of our experiments when these controls are used with DP based optimizers.

The higher revenues obtained by DP models relative to EMSRb with full PODS passenger choice model relative to the First Choice Only Choice model motivated the discussions of competitive feedback effects like inter-airline spill-in and passenger sell-up which help DP outperform standard EMSRb (or DAVN). We observed that due to its inherent aggressive seat protection for higher fare classes, DP causes more passengers to sell-up to higher fare classes which translates to an increase in revenues. However, when the competitor switches to a more sophisticated RM system (in our experiments, we used symmetric RM systems), we observe that the revenues obtained by DP decreased in the majority of experiments indicating that most of the gains obtained by using DP when Airline 2 uses standard EMSRb are due to its inherent aggressive seat protection policy for higher classes which benefits from inter-airline spill in and passenger sell-up within the airline using DP. We also discussed that though the aggressiveness of individual components in an RM system may benefit the airline, being intelligently aggressive is the key to a good RM system. We demonstrated through the examples of DAVN-DP with hybrid forecasting or Fare Adjustment with UDP and hybrid forecasting that these RM systems become overly aggressive and lose revenues relative to DAVN.

We observed in the single flight leg simulation experiment that a very high demand factor results in DP generating higher revenues than EMSRb. In PODS simulations, we observed that DP outperforms EMSRb if there is more higher class demand compared to lower class demand. A practical example of such a situation is when airlines add restrictions in their fare structure. We simulated the change from a semi-restricted fare structure to a fully restricted fare structure. This resulted in an increased demand for higher fare classes and decreased demand for lower fare classes. Under these structural conditions where there is more demand for higher classes relative to lower classes, DP benefits more relative to EMSRb due to its aggressive seat protection policy for higher classes. In our experiments, we observed that DP outperformed EMSRb by a bigger margin in a fully restricted network as compared to a semi-restricted network.

To summarize, although DP did not outperform EMSRb in the single flight leg simulations performed in Chapter 3, our simulation results from PODS in Chapter 5 show that DP can perform better than EMSRb (or standard DAVN) due to its inherent aggressiveness in seat protection for higher classes, rather than its theoretical optimization advantages. In a competitive environment, DP's aggressiveness helps DP gain more revenues due to competitive feedback effects like inter-airline passenger spill-in, and passenger sell-up within the airline.

6.3 Directions for Future Research

In this section, we suggest possible directions for future research activities. We first discuss two of the experimental extensions to our work and then suggest two new theoretical enhancements that can be made to enhance the performance of DP methods.

6.3.1 Experimental Extensions

Validation of Results in Bigger Networks

In this thesis, we discussed results in a competitive network with semi-restricted and fully restricted fare structures consisting of two competing airlines with overlapping single market and hub-and-spoke networks. We believe that airlines typically offer fare structures comprising of both restricted and semi restricted fare structures in the real world. Larger networks have been developed by the PODS consortium and they consist of 4 airlines with different sizes and markets, in which one of the airlines represents a low-cost carrier that offers a fare structure with more compressed fares and fewer restrictions. We have discussed in this thesis that competitive feedback effects like inter-airline passenger spill and passenger sell-up within the airline increase the revenues generated by DP based methods due to DP's aggressive seat protection policy for higher fare classes. We believe that it would be interesting to observe the performance of DP in a more complex and asymmetrical network environment which would be a more accurate representation of the domestic airline network. It might be expected that the competitive feedback effects will be lower in these networks, thereby, adversely affecting DP's revenue performance.

Sell-up Estimation

We have discussed that passenger sell-up is an important reason why DP methods work in our simulation experiments. In this thesis, we have used standard input "FRAT5C" as a measure of passengers' willingness to sell up to a higher fare class when used with Hybrid Forecasting and Fare Adjustment. As a first step, it would be interesting to see the performance of DP methods with other standard FRAT5 curves, namely "FRAT5A" and "FRAT5E", which represent different levels of passenger sell-up. We have already reported that intelligent aggressiveness really pays off in a competitive environment. Hence, the sensitivity of the DP performance on these assumed sell-up rates should be studied.

It is also important that some work be done in estimating passengers' sell up as best as possible. PODS consortium is actively developing new approaches to model this problem and we expect that DP methods would benefit a great deal with accurate estimates of passenger sell-up. Tam (2008) discussed the use of these methods with DP models in unrestricted fare structures but he concluded by saying that "these estimators are far from proven products". We could start off by using the existing techniques to model passenger sell-up and use them in conjunction with DP models to study the effects on the performance of DP models.

6.3.2 Theoretical Extensions

Incorporating Forecast Variance in DP Models

We discussed in Chapter 3 that the benefits of having a dynamic booking limit policy (as in DP) are offset by the fact that the assumed demand variance in DP models is not equal to the realized variance (DP models assume Poisson variance). An important question from theoretical standpoint is how to incorporate the forecast demand variance in DP models. The study done by Walczak (2006) can be a good starting point in this direction.

Approximate Dynamic Programming

We discussed the DP models from the perspective of revenue gains. We did not discuss the computational properties of DP models relative to a heuristic like EMSRb which is extremely simple to implement. The simulation environment in our thesis had just 6 fare classes whereas in real world, airlines can have up to 26 fare classes. As the prolem size grows bigger, the state-space in DP explodes from a computational perspective and DP might become unusable due to "curse of dimensionality". A simple yet effective way to work around this problem is using Approximate Dynamic Programming (ADP). An added benefit of using ADP is that it eliminates the need to know the transition matrix/kernel explicitly. It uses forward recursion by sampling states and consists of a big Monte Carlo simulation as the wrapper. As we are not visiting each and every state (like in DP), efficient sampling techniques for state sampling can reduce the computational time by a big margin. The value function can also be approximated, for instance by piecewise linear functions, to make the deterministic linear optimization at each step computationally tractable. We believe that ADP will have to lead the way in future research in this area if DP methods are to be made applicable in real world airline Revenue Management Systems.

Bibliography

C. Barnhart, P. P. Belobaba, and A. R. Odoni. Applications of operations research in the air transport industry. *Transportation Science*, 37(4):368–391, November 2003.

P. P. Belobaba. Airline yield management: An overview of seat inventory control. *Transportation Science*, 21(2):63–73, May 1987a.

P. P. Belobaba. *Air Travel Demand and Airline Seat Inventory Management*. Phd. thesis, Massachusetts Institute of Technology, Flight Transportation, 1987b.

P. P. Belobaba. Application of a probabilistic decision model to airline seat inventory control. *Operations Research*, 37(2):183–197, March-April 1989.

P. P. Belobaba. Optimal vs. heuristic methods for nested seat allocations. In *Proc. of AGIFORS Reservations and Yield Management Study Group*, pages 28–53, Brussels, Belgium, 1992.

P. P. Belobaba. *Airline Differential Pricing For Effective Yield Management*, chapter 27, pages 349–361. The Handbook of Airline Economics. The Aviation Weekly Group of the McGraw-Hill Companies, New York, NY, 1995.

P. P. Belobaba. *Airline Network Revenue Management: Recent Developments and State of the Practice*, chapter 10, pages 141–156. The Handbook of Airline Economics. The Aviation Weekly Group of the McGraw-Hill Companies, New York, NY, 2002.

P. P. Belobaba and C. Hopperstad. Alternative rm algorithms for unrestricted fare structures. In *AGIFORS Reservations and Yield Management Study Group Meeting*, Auckland, New Zealand, March 2004.

P. P. Belobaba and L. R. Weatherford. Comparing decision rules that incorporate customer diversion in perishable asset revenue management situations. *Decision Sciences*, 27(2):343–363, Spring 1996.

P. P. Belobaba, A. Odoni, and Cynthia Barnhart. *The Global Airline Industry*, chapter 4, page 83. John Wiley and Sons Ltd., 2009.

D. Bertsimas and S. de Boer. Simulation-based booking limits for airline revenue management. *Operations Research*, 53(1):90–106, January-February 2005.

D. Bertsimas and I. Popescu. Revenue management in a dynamic network environment. *Transportation Science*, 37(3):257–277, aug 2003.

G. R. Bitran and S. V. Mondschein. An application of yield management to the hotel industry considering multiple day stays. *Operations Research*, 43(3):427–443, May-June 1995.

E. A. Boyd and R. Kallesen. The science of revenue management when passengers purchase the lowest available fare. *Journal of Pricing and Revenue Management*, 3(2):171, 2004.

S. J-C. Bratu. Network value concept in airline revenue management. Master's thesis, Massachusetts Institute of Technology, Dept. of Aeronautics and Astronautics, 1998.

S. L. Brumelle and J. I. McGill. Airline seat allocation with multiple nested fare classes. *Operations Research*, 41(1):127–137, January-February 1993.

J. Buhr. Optimal sales limits for 2-sector flights. In *Proc. of the Twenty Second Annual AGIFORS Symposium*, pages 291–303, Lagonissi, Greece, October 1982.

E. Carrier. Modeling airline passenger choice: Passenger preference for schedule in passenger origin destination simulator (pods). Master's thesis, Massachusetts Institute of Technology, Department of Civil and Environmental Engineering, 2003.

R. L. Cléaz-Savoyen. Airline revenue management methods for less restricted fare structures. Master's thesis, Massachusetts Institute of Technology, Dept. of Aeronautics and Astronautics, 2005.

R. E. Curry. Optimal airline seat allocation with fare classes nested by origins and destinations. *Transportation Science*, 24(3):193–204, August 1990.

Y. Feng and B. Xiao. A dynamic airline seat inventory control model and its optimal policy. *Operations Research*, 49(6):938–949, November-December 2001.

T. Fiig, K. Isler, C. Hopperstad, and P. Belobaba. Optimization of mixed fare structures: Theory and applications. *Journal of Revenue and Pricing Management*, 9 (1–2):152–170, Jan 2010.

G. Gallego and G. Van Ryzin. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40(8):999–1020, aug 1994.

G. Gallego and G. Van Ryzin. A multiproduct dynamic pricing problem and its applications to network yield management. *Operations Research*, 45(1):24–41, January-February 1997.

T. Gorin. Airline revenue management: Sell-up and forecasting algorithms. Master's thesis, Massachusetts Institute of Technology, Dept. of Aeronautics and Astronautics, 2000.

J. Grandell. *Mixed Poisson Processes*, chapter 2, page 13. Chapman and Hall, 1997.

J. C. Guo. Estimation of sell-up potential in airline revenue management systems. Master's thesis, Massachusetts Institute of Technology, Sloan School of Management, 2008.

Craig Hopperstad. *Passenger O/D Simulator PODS*. Hopperstad Consulting, Seattle, v.10 edition, September 2005.

A. Kimms and M. Müller-Bungart. Simulation of stochastic demand data streams for network revenue management problems. *OR Spectrum*, 29(1):5–20, jan 2007.

M. E. Kuhl and J. R. Wilson. Advances in modeling and simulation of nonstationary arrival processes. In *Proc. of the 2009 INFORMS Simulation Society Research Workshop*, pages 1–5, Coventry, United Kingdom, June 2009.

C. J. Lautenbacher and S. Stidham Jr. Airline yield management with overbooking, cancellations, and no-shows. *Transportation Science*, 33(2):147–167, may 1999.

T. C. Lee and M. Hersh. A model for dynamic airline seat inventory control with multiple seat bookings. *Transportation Science*, 27(3):252–265, aug 1993.

P. A. W. Lewis and G. S. Shedler. Simulation of non-homogenous poisson process by thinning. *Naval Research Logistics Quarterly*, 26:403–413, 1979.

E. L'Heureux. A new twist in forecasting short-term passenger pick-up. In *Proc. of the Twenty Sixth Annual AGIFORS Symposium*, Bowness-on-Windermere, UK, 1986.

K. Littlewood. Forecasting and control of passenger bookings. In *Proc. of the Twelfth Annual AGIFORS Symposium*, pages 95–117, Nathanya, Israel, October 1972.

J. I. McGill and G. J. van Ryzin. Revenue management: Research overview and prospects. *Transportation Science*, 33(2):233–256, May 1999.

M. H. Reyes. Hybrid forecasting for airline revenue management in semi-restricted fare structures. Master's thesis, Massachusetts Institute of Technology, Dept. of Civil and Environmental Engineering, 2005.

H. Richter. The differential revenue method to determine optimal seat allotments by fare type. In *Proc. of the Twenty Second Annual AGIFORS Symposium*, pages 339–362, Lagonissi, Greece, October 1982.

L. W. Robinson. Optimal and approximate control policies for airline booking with sequential nonmonotonic fare classes. *Operations Research*, 43(2):252–263, March-April 1995.

M. Rothstein. An airline overbooking model. *Transportation Science*, 5(2):180–192, May 1971.

M. Rothstein. Airline overbooking: Fresh approaches needed. *Transportation Science*, 9(2):169–173, May 1975.

R. W. Simpson. *Using Network Flow Techniques to Find Shadow Prices for Market and Seat Inventory Control.* Memorandum m89-1, Massachusetts Institute of Technology, Flight Transportation Laboratory, 1989.

B. C. Smith and C. W. Penn. Analysis of alternate origin-destination control strategies. In *Proc. of the Twenty Eighth Annual AGIFORS Symposium*, pages 123–144, New Seabury, Massachusetts, United States, October 1988.

J. Subramanian, S. Stidham Jr, and C. J. Lautenbacher. Airline yield management with overbooking, cancellations, and no-shows. *Transportation Science*, 33(2):136–146, may 1999.

K. T. Talluri and G. J. Van Ryzin. *The Theory and Practice of Revenue Management*, volume 68 of *International Series in Operations Research & Management Science.* Kluwer Academic Publishers, 2004.

C. F. W. Tam. Airline revenue management based on dynamic programming incorporating passenger sell-up behavior. Master's thesis, Massachusetts Institute of Technology, Department of Civil and Environmental Engineering, 2008.

T. Vanhaverbeke. Revenue management based on dynamic programming in unrestricted and simplified fare structures. Master's thesis, Massachusetts Institute of Technology, Department of Civil and Environmental Engineering, 2006.

B. Vinod. *Origin and Destination Yield Management*, chapter 27, pages 459–468. The Handbook of Airline Economics. The Aviation Weekly Group of the McGraw-Hill Companies, New York, NY, 1995.

D. Walczak. Modeling high demand variance in dynamic programming. *Journal of Revenue and Pricing Management*, 5(2):94–101, Jul 2006.

K. Wang. Optimal seat allocation for multi-leg flights with multiple fare types. In *Proc. of the Twenty Third Annual AGIFORS Symposium*, pages 225–237, United States, October 1982.

L. R. Weatherford. *Perishable Asset Revenue Management in General Business Situations.* Phd. thesis, University of Virginia, Darden Graduate School of Business Administration, 1991.

L. R. Weatherford, S. E. Bodily, and P. E. Pfeifer. Modeling the customer arrival process and comparing decision rules in perishable asset revenue management situations. *Transportation Science*, 27(3):239–251, aug 1993.

E. L. Williamson. Comparison of optimization techniques for origin-destination seat inventory control. Master's thesis, Massachusetts Institute of Technology, Flight Transportation Laboratory, 1988.

E. L. Williamson. *Airline Network Seat Inventory Control: Methodologies and Revenue Impacts*. Phd. thesis, Massachusetts Institute of Technology, Flight Transportation, 1992.

R. D. Wollmer. An airline seat management model for a single leg route when lower fare classes book first. *Operations Research*, 40(1):26–37, February 1992.

R. Wysong. A simplified method for including network effects in capacity control. In *Proc. of the Twenty Eighth Annual AGIFORS Symposium*, pages 113–121, New Seabury, Massachusetts, United States, October 1988.

W. Zhao and Y. Zheng. Optimal dynamic pricing for perishable assets with nonhomogenous demand. *Management Science*, 46(3):375–388, mar 2000.

W. Zhao and Y. Zheng. A dynamic model for airline seat allocation with passenger diversion and no-shows. *Transportation Science*, 35(1):80–98, feb 2001.

J. S. Zickus. Forecasting for airline network revenue management: Revenues and competitive impacts. Master's thesis, Massachusetts Institute of Technology, Civil and Environmental Engineering, 1998.