Airline Passenger Cancellations: Modeling, Forecasting and Impacts on Revenue Management

By
Oren Petraru
B.A. Economics and Business Administration,
Hebrew University of Jerusalem, Israel (2013)

Submitted to the Department of Civil and Environmental Engineering in partial fulfillment of the requirements of the degree of

Master of Science in Transportation

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2016

Signature Redacted

Author ....
Department of Civil and Environmental Engineering
May 13, 2016

Signature Redacted

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

Signature Redacted

Accepted by ....
Heidi Nepf
Donald and Martha Harleman Professor of Civil and Environmental Engineering
Chair, Graduate Program Committee

Airline Passenger Cancellations: Modeling, Forecasting and Impacts on Revenue Management

By

Oren Petraru

Submitted to the Department of Civil and Environmental Engineering
on May 13, 2016 in partial fulfillment of the requirements
of the degree of Master of Science in Transportation

ABSTRACT

Passenger demand forecasting, and subsequently passenger cancellation forecasting, are important components in any airline revenue management (RM) system. Passenger cancellations can potentially lead to flights leaving with empty seats and thus to loss of revenues. Airlines need accurate cancellation forecasting tools in order to properly compensate for cancellations, or in other words, overbook flights above their physical capacity. At the same time, airlines need to be cautious not to overbook too aggressively. If a flight is still overbooked at time of departure, not all passengers are able to board and those left behind need to be compensated and re-accommodated.

This thesis focuses on modelling and forecasting passenger cancellations using the PODS booking simulation tool. Several methods for cancellation forecasting and overbooking are presented and their impacts are tested under different demand, competition and RM strategy settings. All methods are based on time series modeling of historical observations. However, the methods differ in terms of the data they use and the cancelled bookings they compensate for. The potential contribution of Passenger Name Record data (PNR) to more accurate cancellation forecasting is discussed as well.

Simulation results indicate that the ticket revenue gains due to cancellation forecasting and overbooking range between 1.15% and 4.16%, depending on the cancellation forecasting method used and the level of overbooking aggressiveness. However, aggressive overbooking increases the negative effect on revenues due to the costs associated with denied boardings. Therefore, after taking into account these costs, the net revenue gains range between 0.06% and 2.79%. For airlines with high cancellation rates, the magnitude of the gains from cancellation forecasting and overbooking is even greater, reaching 3.59% in net revenue improvements.

Thesis Supervisor: Peter P. Belobaba

Title: Principal Research Scientist of Aeronautics and Astronautics

Acknowledgements

First and foremost, I would like to start by thanking my advisor, Dr. Peter Belobaba, for his guidance and support throughout my time at MIT. It was truly an honor and privilege to take his classes and be a research assistant for the PODS research consortium. I really enjoyed the opportunity to learn from his tremendous knowledge and experience about the airline industry and revenue management. I highly appreciate his involvement and "hands on" approach in my work.

Second, I would like to thank Craig Hopperstad, the original developer of the Passenger Origin Destination Simulator (PODS) for contributing his knowledge about the mechanism behind PODS with patience. His work provided a crucial basis for understanding my research.

Third, I would like to thank Dr. Nicole Adler from the Hebrew University's School of Business. I really appreciate the opportunity she gave me at the beginning of my third year in undergrad to be her research assistant and learn so much about commercial air transportation. Working with her opened a door to my graduate studies at MIT, and for that I will be forever grateful.

Fourth, I would like to thank "Team PB" (in alphabetical order): Adam, Alex B., Alex S., Daniel, Germán, Matthieu and Mike. I consider myself extremely lucky to be a part of such a smart, talented and fun group of people. It was a true pleasure to spend time with you and learn from you all, either in the lab or in a "PODS destination". I will really miss our long conversations about airlines, travel and life in general. Thank you for being a part of my life.

Fifth, I would like to thank the ICAT lab mates: Sathya, Tamas, Henry, Mayara and Morrisa. Thank you all for making my time at MIT a fun one, especially at stressful times. It was great to know you all.

Sixth, I would like to thank the RM/OR team at American Airlines: Marcial, Qiuting, Michael, Gonzalo, Yuxi, Soufiane. I really appreciate the opportunity to be exposed to the fascinating work place dynamics and the challenges of a "real world" airline as a summer intern. I learned a lot during my time there, and I am grateful for their contribution to my thesis.

Seventh, I would like to thank my family and friends outside MIT. I really appreciate all your support and faith in me during my time at MIT but also before I got in. It is always good to know that no matter what, I am surrounded by extraordinary people who will always be there for me.

Last, but definitely not least, I want to thank the love of my life, my wife Emily. I would not have made it through the last two years without your unconditional support, especially in difficult and stressful times. Thank you for believing in me so much, for your wise words and for putting up with me no matter what.

Table of Contents

List of Figures

Chapter 1: Introduction

The concept of airline revenue management (RM) has grown after the deregulation of the US airline industry in the late 1970s. While before deregulation airlines had to set ticket prices according to the mileage-based formula used by the Civil Aeronautics Board and charge all passengers the same price, after deregulation airlines introduced different techniques such as price discrimination, product differentiation and fare restrictions. These techniques are based on the notion that passengers have different levels of willingness-to-pay (WTP) and different characteristics in terms of trip purpose and sensitivity to price and time. Airlines attempt to maximize their revenues by introducing multiple booking classes with different fares and restrictions such as Saturday minimum stay and advance purchase to ensure that passengers with high WTP purchase high fare tickets while passengers with lower WTP purchase lower fare tickets to fill up the remaining unoccupied seats (Belobaba, 2016). The RM systems allocate seats to each fare product based on the demand forecasts generated for each future flight. Low fare products will have booking limits in order to prevent passengers with high WTP who are not limited by restrictions from purchasing tickets with fares lower than their WTP ("buying down"). High fare products will have protection levels in order to ensure that requests for high fare seats will never be turned down.

1.1 Importance of Demand and Cancellation Forecasting

One of the core elements in an airline's RM system is its ability to forecast demand for any given flight as accurately as possible. After the forecast is produced, usually based on historical bookings, the RM system then determines the seat allocation that will maximize revenue for the airline. However, just as in any other business, forecasting demand for air travel is a challenging task for airlines as demand can be highly volatile due to several factors such as seasonality, holidays, economic shifts, geo-political disputes, epidemic outbursts and terror attacks to name a few. Even with the advancement in computational applications which allow airlines to produce many different forecasts in a very short time (Neuling et al., 2003) forecasts are rarely a precise prediction of future demand.

Inaccuracy of demand forecast may lead to unfavorable consequences to the airline. If demand forecasts are higher than the actual demand level, the RM system may allocate fewer seats to low WTP passengers as it expects demand to be sufficient enough to fill the remaining seats with high WTP passengers. However, since there is not enough actual demand, flights may leave with empty seats to their destination. Since an airline seat on a given flight and date is a perishable product/service that cannot be recuperated after the flight leaves for its destination, the airline loses potential revenue. If demand forecasts are lower than actual demand level, the RM system may allocate too many seats for low WTP passengers as it predicts there is not enough demand to fill all seats. In this case, the RM system does not protect enough seats for passengers with high WTP who book later in the booking process. In both cases, airlines do not maximize their revenues as they do not optimize the number of seats made available to each fare class.

One of the main issues airlines have to take into consideration while building models for demand forecasts is the phenomenon of passenger cancellations. The fact that passengers can purchase their tickets well in advance before a departure date increases the probability that at some point in time until departure, passenger plans will change and hence will force them to either cancel their itinerary completely, or reschedule their trips to different dates. Any cancellation by a passenger in the booking process adds disturbance to the demand forecasting system and can result in revenue loss. The potential loss of revenue also increases the later the passenger cancels his/her itinerary as airlines have less time and flexibility to try and compensate for the revenue loss. A sub-category of cancellation behavior is called "no-show" and this behavior refers to cases where booked passengers fail to show up at the gate for departure. In these cases passengers do not cancel their itineraries prior to departure and therefore airlines do not have the time to try and sell the seat to another passenger and loss of revenue is almost guaranteed unless anticipated by the airline.

Forecasting cancellations could be an easy task if all passengers were the same. However the reality is that passengers differ from one another in myriad ways and so is their likelihood of cancelling their trips or not showing up to their flight on day of departure. Belobaba et al. (2016) discuss two passenger types with their unique characteristics: "leisure" and "business" passengers. "Business" passengers are described as insensitive to price and time sensitive who go on business trips of up to a few days during the week days, hence weekends and holidays are excluded, and book their trips closer to departure. "Leisure" passengers are characterized as sensitive to price and insensitive to time who travel for the purpose of vacation or visiting friends and family and therefore travel for at least a few days, including weekends and holidays, and book their trips far in advance. In order to ensure each type of passenger is paying according to its WTP, the airlines use these characteristics to determine the restrictions associated with each fare class.

Airlines implement several methods of attempting reduce the risk of revenue loss associated with cancellations. One method is applying "non-refundable" or "non-changeable" restrictions on some fare class. These restrictions can reduce the probability of passengers cancelling after they had already booked their trips and at the same time incentivize passengers to book their trips only if their plans are firm. For "business" passengers these restrictions may be more inconvenient as their plans might change even after booking and hence airlines usually allow bookings to be refundable or changeable in higher fare classes. Allowing a passenger to get a refund or change their itineraries in exchange for paying a higher fare for their booking reduces the revenue loss expected in case of a cancellation.

Another method airlines implement to reduce losses associated with passenger cancellations or "no-show" is overbooking. By overbooking, airlines allow bookings to exceed the physical seat capacity on a given flight and thus reduce the probability of flights leaving with empty seats. Though, just as in the case of demand forecasting, the cancellation forecast needs to be as accurate as possible in order to enable airlines to overbook just enough to fill all seats without exceeding the actual capacity. Airlines have developed cost based models to compute optimal levels of authorized capacity (AU) on flights (Belobaba et al., 2016). The objective of the models is to find the AU where total costs of denied boarding (DB) and spoiled seats are at minimum. The main challenge for the airlines with these models is to produce accurate cost estimates of denied boarding and spoiled seats as some costs are difficult to quantify. It is then up to the airline to decide on a strategy that will be best for the airline economically.

1.2 Motivation for Research

Cancellation forecasting has been a research topic for airlines since the evolution of RM in the early 1980s (Neuling et al., 2003). Accurate cancellation forecasting is a crucial element for an airline's profitability and much effort has been invested by the airlines in developing models to address this issue. Cancellation rates can be different for each airline as the rates are mainly a factor of the airline's booking restrictions policy and passenger mix. According to Illiescu et al. (2008), cancellation rates reaching 30% or more are not uncommon, hence the importance of accurate cancellation forecasting is understandable.

The most common approach to cancellation forecasting has been, and still is, time series modeling based on historical observations. These models attempt to forecast two types of cancellation rates: cancellation rates for bookings that have already been accepted and cancellation rates for bookings that are yet to be made as the RM system needs these rates for estimates of demand for a flight. The models take in consideration the differences in rates due to specific flight attributes such as flight time, day of week, seasonality and holidays to name a few (Garrow et al., 2004). A more recent approach for cancellation rates estimation has been the use of disaggregate Passenger Name Record (PNR) information from the airlines' reservation systems. PNR data contains more detailed information about each booking which can contribute to better estimation of cancellation rates on an aggregate level as the probability of cancellation of an individual booking is not required by the RM system (Romero Morales et al., 2010).

Despite the research done so far on the topic of passenger cancellation forecasting, very little research has been done on the impact of different cancellation forecasting and overbooking methods on airlines' RM systems, taking account of passenger choice. Specifically, the interaction between cancellation forecasting and overbooking methods with different RM systems, forecasters, under different demand levels and cancellation rates and the impact on airlines' performance has yet to be investigated. Therefore this thesis aims to analyze the consequences of cancellation behavior as well as different cancellation forecasting and overbooking methods on airline revenues. The Passenger Origin Destination Simulator (PODS) which simulates airline competitive environment and passenger choice behavior is used in this thesis. The passenger choice simulation has been expanded to include cancellation behavior for the purpose of this thesis. The aspects to be investigated include the cancellation forecasting and overbooking method used by the airline, the cancellation rate levels, the RM system used, the type of forecast fed into the optimizer and the RM capabilities of the competitors. Each scenario will present the results of the simulation runs with the main focus being on the revenues, load factor and yield of one airline in a competitive environment.

The thesis begins with this introduction and proceeds to Chapter 2 which includes a literature review regarding cancellation forecasting methods and passenger cancellation behavior in the airline industry. Chapter 3 provides an overview of PODS and presents the adjustments made (and the assumptions behind them) in order to incorporate cancellation behavior in the simulator. The chapter includes several methods for passenger cancellation forecasting based on time-series analysis of historical observations. The cancellation rates that are calculated by each method will be used towards overbooking of remaining capacity in order to compensate for the potential loss of revenues.

Chapter 4 presents the results of a series of test simulations done in PODS. The first couple of tests will compare the results of one airline in a competitive environment using different cancellation forecasting methods with different levels of overbooking aggressiveness. One test will compare the results for an airline using different cancellation rate aggregations. Another test will impact of a different overbooking approach on the performance measures. The rest of the tests will show the results under different competitive, cancellation rates and RM settings.

Chapter 5 will discuss the importance of using detailed passenger data that is included in every airline's databases towards the development of more sophisticated and thus more accurate cancellation forecasting models. The chapter will give some examples of PNR attributes that could be used in better understanding passenger cancellation behavior. Chapter 6 will conclude this thesis and briefly discuss ideas for future research.

Chapter 2: Literature Review

The subject of passenger cancellation forecasting in airline RM systems has been covered by the academic literature for several decades. In McGill and van Ryzin's (1999) extensive review of 40 years of literature on transportation RM, it is claimed that research on reservations control and overbooking dates back to the late 1950s. Overbooking calculations were based on forecasting of probability distributions of the number of passengers who showed up for their flights, and therefore overbooking research motivated research on forecasting of passenger cancellations and no-shows. Another review on the topic of overbooking can be found in Rothstein (1985), where the motivation for overbooking, the advantages and early practices of overbooking are discussed.

More recent literature regarding cancellation and overbooking in RM systems can be found Gosavi et al. (2005) where an integrated simulation based approach is developed for solving seat allocation problems in the airline industry taking into consideration passenger cancellations and overbooking of flights. Subramanian et al. (1999) developed a Markov decision process model for airline seat allocation for single leg flights with multiple fare classes allowing for cancellations, no-shows and overbooking. Both examples do not include customer choice behavior. Sierag et al. (2015) introduce a customer choice cancellation model as an extension to RM model proposed by Talluri and van Ryzin (2004). The customer choice model is based on a Markov decision process and dynamic programming formulations. The authors conclude that failure to take into cancellations in RM models can lead to a revenue loss of up to 20 percent.

The importance of cancellation forecasting accuracy to airlines' revenue maximization is often discussed in the literature. Lawrence et al. (2003) assert that underestimation of no-shows leads to loss of potential revenues due to unoccupied seats, while overestimation can result in a cost penalty for the airlines as they need to compensate for denied boardings. Romero-Morales et al. (2010) add that in revenue management system using dynamic pricing overestimation of cancellation rates can result in underestimation of demand by the system and hence prices will be set too low in order to attract more demand. Chatterjee (2001) demonstrates the necessity in forecasting cancellation rates for bookings-in-hand (BIH) and booking-to-come (BTC) instead of relying solely on forecast of net demand due to the uncertainty (i.e. variance) of BIH and BTC. If BTC are not taken into account, there is a risk of underestimating the total variance of the net demand.

Most of the work done on passenger cancellation forecasting methods is focused on two main approaches. The first approach is traditional time series forecasting. This approach is a relatively straightforward and intuitive approach which is being used in many fields besides transportation. According to Garrow and Koppelman (2004), airlines forecast no-shows rates using time-series models based on historical class or cabin no-show rates. Lemke et al. (2008, 2009) investigate several methods for forecasting using time series data. In addition to the commonly used methods of simple averaging and exponential smoothing, other methods such as regressions, decomposition and Theta models, autoregressive integrated moving average models and non-linear forecasting are also proposed as methods for forecasting in airline RM systems. Airlines forecast cancellation rates on fare class, day of week, point of sale and OD levels and not on specific passengers' characteristics levels such as business or leisure. For this reason, Garrow and Koppelman (2004) argue that forecasting models based on historical bookings cannot make accurate predictions when passenger and/or itinerary mix changes.

The second approach for cancellation forecasting that has gradually become more popular among airlines in recent years is Passenger Name Record (PNR) based forecasting. In this approach, the airline uses the attributes that are included in each PNR to develop a forecasting model at a more disaggregate level. Flight times and dates, origin and destination (OD) airports, seasonal aspects and passenger attributes such as frequent flyer program membership and number of passengers travelling together are just some of the attributes that can be retrieved from the PNR data and be used in models for cancellation forecasting. The data can be used in models for cancellation or no show predictions either at a passenger level or cabin level as suggested by Lawrence et al. (2003). Passenger level models can be implemented by using decision trees, Naïve Bayes algorithms and other probabilistic models while cabin level models can implemented by using linear regressions or probabilistic models as well. Garrow and Koppelman (2004) suggest a multinomial logit model that predicts whether or not each passenger (i.e. PNR) will or will not show up for a flight. Neuling et al. (2003) also emphasize the fact the due to the amount of information required for successful implementation of PNR based forecasting models, several steps need to be taken in such as understanding the information included in the PNR data and selecting proper attributes to be used later on in the predictive models. Failure to do these basic yet crucial steps may result in highly biased results.

The superiority of the PNR based model over the time series model is mentioned in several papers. For example, according to Lawrence et al. (2003) a cabin level passenger based model can produce between 0.4 to 3.2 percent revenue gains over the conventional model. Nevertheless, the literature acknowledges the shortcomings of the PNR model. Airlines forecast cancellation rates by time periods before departure and the closer the time period is to departure the higher number of total bookings it contains. Since PNR based models are dependent on number of bookings, they perform better in forecasting cancellation the closer the bookings are to departure date. PNR based no show models are incapable of producing a complete forecast early in the booking process when few bookings are actually available. Furthermore Romero Morales et al. (2010) address the complex dynamics in the behavior of passengers in the different stages of the booking horizon and point out that the set of relevant attributes is very diverse in different stages of the booking horizon.

In addition to the shortcomings of each model individually, there is also a general consensus that using a single model can be risky as it can only be accurate to a certain degree due to the changes in data and performance over time. Combined models can complement each other and increase the accuracy of the forecast due to the improved ability to adapt to changes in the data. Lemke et al. (2008, 2009) discuss linear, non-linear as well as adaptive forecast combinations for improved accuracy of models. They emphasize the fact that combining forecasting does not always improve performance however generation of forecasts on different levels of aggregation of data can result in a significance performance gain. Lawrence et al. (2003) suggest producing a weighted average of historical and passenger based forecasts with increased weights assigned to the passenger based forecast as more booking are realized. Romero Morales et al. (2010) propose building multiple models for different stages of the booking horizon or a single model that takes into account time dependency.

Illiescu et al. (2008) used ticketing data from the Airline Reporting Corporation for group ticketing within 90 days before departure and the occurrence of refund and exchange events in a discrete time proportional odds (DTPO) model with a prospective time scale to model airline passengers' cancellation behavior. The authors discuss three main findings. First, in general higher cancellation rates are observed for recently purchased tickets and for tickets whose flight departure dates are near. This finding is in line with "bath-tub" shaped passenger cancellation pattern which is used in PODS and will be described in further detail later in this thesis. Second, passengers travelling in groups have lower probabilities for cancelling in comparison passengers who are travelling alone. Third, based on variables associated with leisure passengers such as Saturday night stay and end of week day of outbound departure, tickets that are booked by leisure passengers have lower cancellation rates.

Belobaba (2016) presents the necessity for airlines to overbook flights as well as different approaches for determining authorized capacity (AU) for flights. The approaches include judgement of human analysts based on their previous experience with passenger cancellation and no-shows, deterministic model which determines AU by assuming the actual no-show rate is known with certainty and probabilistic or risk based model which incorporates the variance in the distribution of no-show rates for future flight departures. Airlines can make adjustments to the probabilistic models based on their tolerance for denied boardings and spoiled seats. The probabilistic model has been extended to another model, the cost-based overbooking model, which not only takes into account the uncertainty in no-show behavior but also accounts for the costs airlines assign to denied boarding passengers spoiled seats and attempts to find the optimal AU where total overall costs are minimized. Calculating the exact costs of denied boarding and spoiled seats is not an easy task as these costs cannot be completely quantified in monetary terms.

Karaesman and van Ryzin (2004) present a two-period optimization model to determine joint overbooking levels for multiple class setting where substitution among classes is allowed. Results show that in some cases accounting for substitution can increase revenues to some extent after accounting for penalties. Based on their findings, there is potential to improve overbooking practices for adjacent flights or multi-cabin flights where substitution options are available. Ignaccolo and Inturri (2000) propose a method to minimize both denied boardings and spoiled seats in flights by monitoring the booking process and using Inference Fuzzy Systems to assist RM analysts with their decisions on AU settings.

The existing literature on passenger cancellation behavior and cancellation forecasting methods raises several points. First, the common practice used by airlines today for cancellation forecasting is time series forecasting based on historical observations. Second, the additional value PNR data can provide for the improvement of the accuracy of cancellation forecasting is recognized. Third, the highest level of accuracy can be provided only if several models are used together. That way each model's deficiency can be compensated by other models. Fourth, there are several approaches for compensating for loss of revenues due to cancellations, and every airline can choose its approach based on its preferences and experience. Fifth, cancellation probabilities vary by time, by type of passengers and by number of passengers in group.

Chapter 3: Overview of PODS and Network U10

All the experiments described in this thesis are based on the Passenger Origin Destination Simulator (PODS). This tool, which was initially developed by Boeing as a model for passenger travel window preferences, now simulates a hypothetical airline competitive environment in which each airline has its own network. Each airline gets to decide on its RM system's components such as the RM optimization model and the forecast method to be used and also on the fare products it offers its customers. Demand for air travel is generated by simulating passengers who wish to travel in specific OD markets and thus need to decide between the various paths, airlines and products available to them. It is important to note that in PODS, just as in reality, airlines have no control on the demand generation process and they do not know if demand actually exists until a booking is made. The airlines must use the various forecasting tools available to them to forecast the number of passengers on each future flight, their arrival process and their cancellation patterns. In the next sections PODS and network U10, the framework for this thesis, will be described in greater detail. The figures in sections 3.1 through 3.3, are from a report by Belobaba (2010) which includes many other charts and figures regarding PODS.

3.1 Overview and Structure

The RM system in PODS consists of three sub-systems: the historical booking database, the forecaster and seat allocation optimizer. Each booking made by a passenger is recorded in the historical booking database. The forecaster then uses the historical booking data together with the current booking levels to forecast bookings for every flight in the future. Based on these forecasts, the optimizer allocates the number of seats made available in each fare class, path and time frame before departure. The availability then defines the passenger choice set from which the passengers have to make a decision, based on their characteristics and associated preferences of schedule times, fares and restrictions on each fare class. The characteristics of each passenger are generated stochastically for each market by the demand generator which is the first component in the passenger choice model. After a booking is made, the RM system once again updates the remaining availability in all fare classes. Figure 1 illustrates the interaction between the Passenger Choice Model and RM systems in PODS during the booking process.

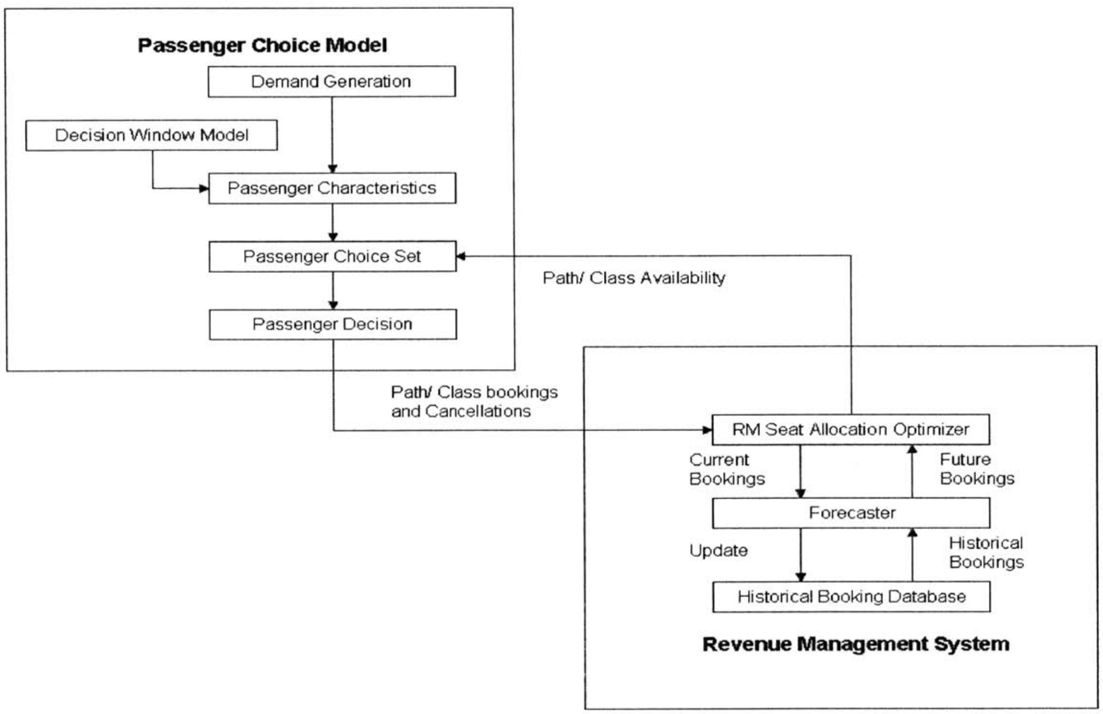

Figure 1: PODS Structure

(Source: Belobaba, 2010)

PODS simulates multiple repetitions of the same departure day and averages the output results. The booking period consists of 63 days, which are divided into 16 time periods or time frames. Figure 2 demonstrates the division of the pre-departure days into their respective time frames. The number of days in each time frame decreases the closer the time frame is to departure. The bookings-in-hand (i.e. the cumulative net bookings) and forecasted bookings-to-come are re-calculated at the beginning of each time frame in PODS.

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

Figure 2: PODS' Time Frames

(Source: Belobaba, 2010)

3.2 Passenger Choice Behavior in PODS

The demand for each market is divided into two passenger types, business and leisure, and each type is randomly assigned a set of characteristics which includes a decision window, a maximum WTP and a set of disutility costs.

1. ***Decision window:*** the time frame within which the passenger is willing to travel. The window is bounded by the earliest departure time and the latest arrival time the passenger will consider. Business passengers have shorter decision windows than leisure passengers due to their time sensitivity and reduced flexibility in travel. If the path and fare class combination is within the decision window it will be included in the passenger choice set and excluded otherwise.
2. ***Maximum WTP:*** the maximum amount the passenger is willing to pay for air travel. It is defined as a function of the input *base fare,* the price which the mean number of passengers in each OD market are willing to pay for travel, and *emult,* the elasticity multiplier such that 50% of passengers are willing to pay *emult*base fare* to travel. The base fare and the *emult* for business passengers are higher than those for leisure passengers due their lower price sensitivity. If a travel option has an out-of-pocket

fare that exceeds the passenger's maximum WTP it will be excluded from the passenger's choice set.

3. ***Set of disutility costs:*** the dollar value of the disutilities arising from the restrictions associated with the fare products in each market and from the re-planning, path quality and unfavorable airline costs. In our case, the cancellation penalty restriction on a ticket increases the disutility for business passengers due to the greater uncertainty in their travel plans compared to leisure passengers. Re-planning cost is the penalty assigned by the passenger if not able to fly within the desired time window. Path quality cost is the penalty assigned by the passenger when flying a connecting (versus a non-stop) itinerary due to the additional time and stress associated with it.

On average, the ratio between business and leisure passengers is 40:60, respectively, which is in line with industry data. Network U10 has four booking curves which describe the arrival pattern of passengers into the booking process based on type of passenger, business and leisure, and type of market, domestic and international, shown in the example in Figure 3. The arrival patterns in the figure differ according to the unique characteristics of each passenger type and are based on industry experience. The leisure arrival curve is above that of the business as leisure passengers tend to book earlier in the booking process than business passengers. The international market arrival curve is above that of the domestic market as international passengers tend to book further in advance due to the additional bureaucracy usually involved in international itineraries such as visa arrangements. At day 21 before departure, 85% of the international leisure passengers, 78% of the domestic leisure passengers, 54% of the international business and 35% of the domestic business passengers have arrived to book.

It should be noted that in reality passenger do not classify themselves as leisure or business when they book and this classification is only used by PODS for the purpose of modeling different types of demand groups. By no means is this supposed to suggest that business passengers will not book a low fare ticket, if the disutility costs associated are not too high, or that a leisure passenger will not book a high fare ticket.

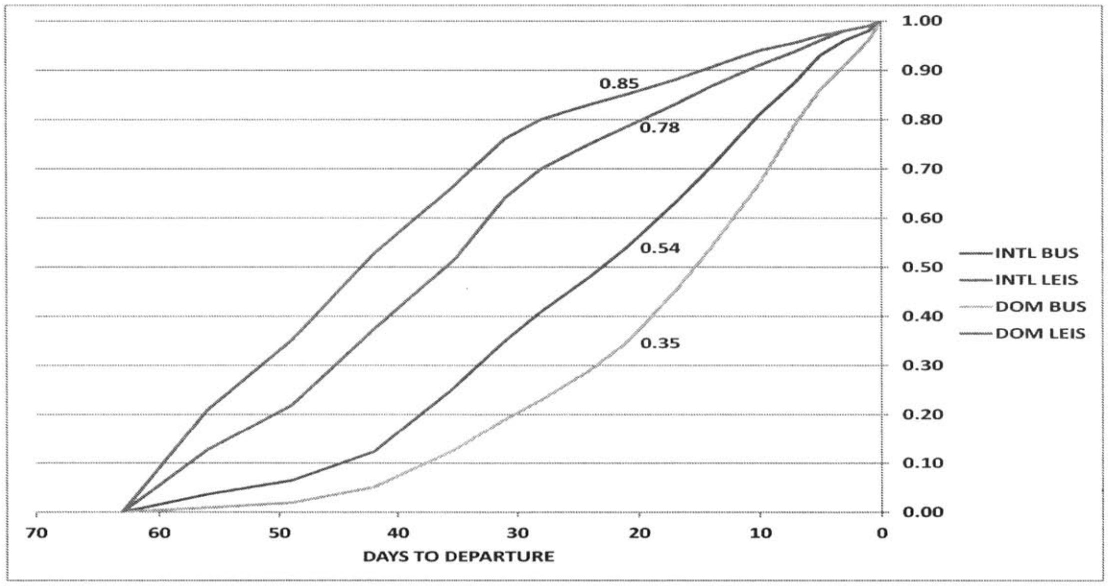

Figure 3: Booking Curves by Passenger and Market Types

(Source: Belobaba, 2010)

3.3 Network U10 Competitive Airline Network

The competitive airline environment used in this thesis consists of four airlines, and each airline is based in its own hub. Airline 1 (AL1) hubs at MSP, Airline 2 (AL2) hubs at ORD, Airline 3 (AL3) hubs at MCI and Airline 4 (AL4) hubs at DFW. Each airline flies to and from its hub to 40 domestic and international spoke cities, 20 east of the hub and 20 west of the hub. The passenger flow goes from West to East. While most markets are served by connecting itineraries, the hubs serve as destination and origin points as well. Overall, the airlines operate 442 legs departure legs per day serving 572 OD markets. Figures 4.1, 4.2, 4.3 and 4.4 illustrate all four airline networks. Airlines 1, 2 and 4 serve both domestic and international markets and represent the "legacy" carriers, while Airline 3 serves domestic markets only and represents the low-cost carrier (LCC).

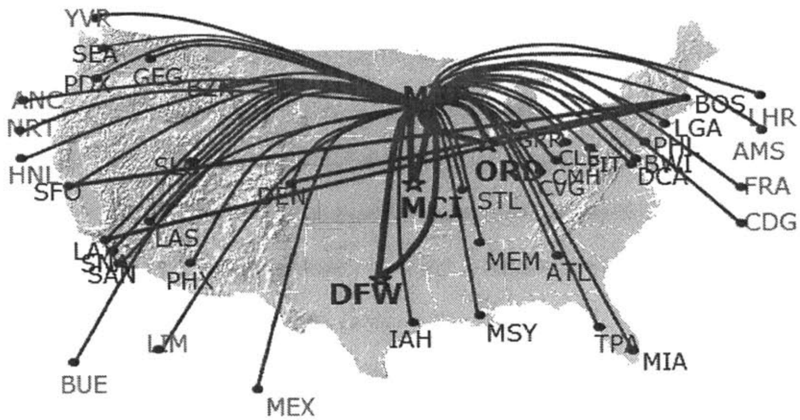

Figure 4.1: AL1 Network- MSP Hub

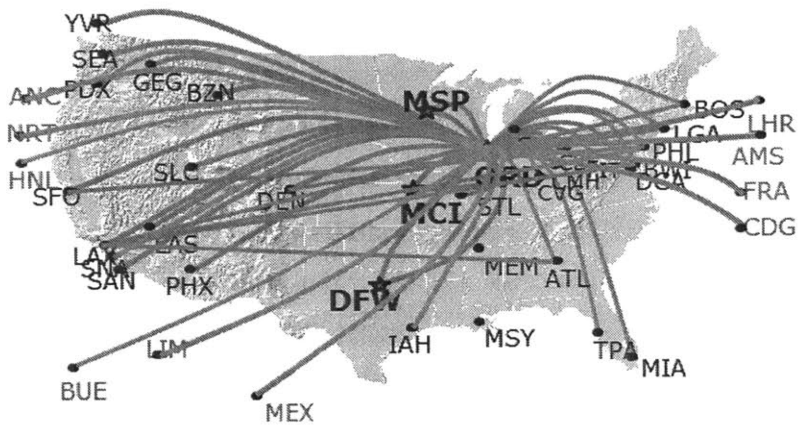

Figure 4.2: AL2 Network- ORD Hub

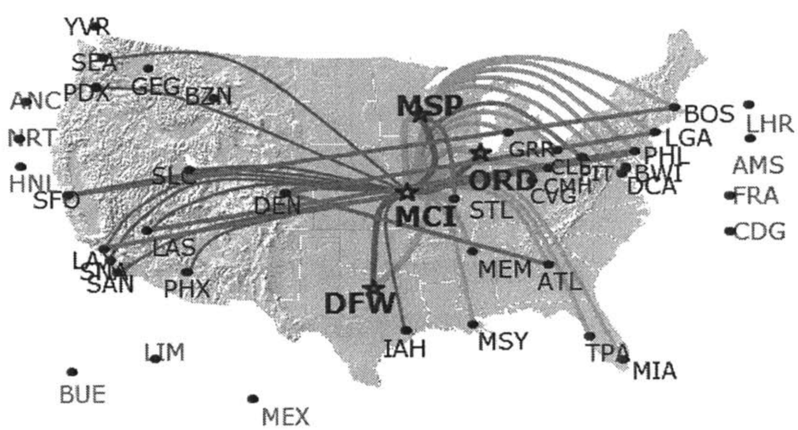

Figure 4.3: AL3 Network- MCI Hub

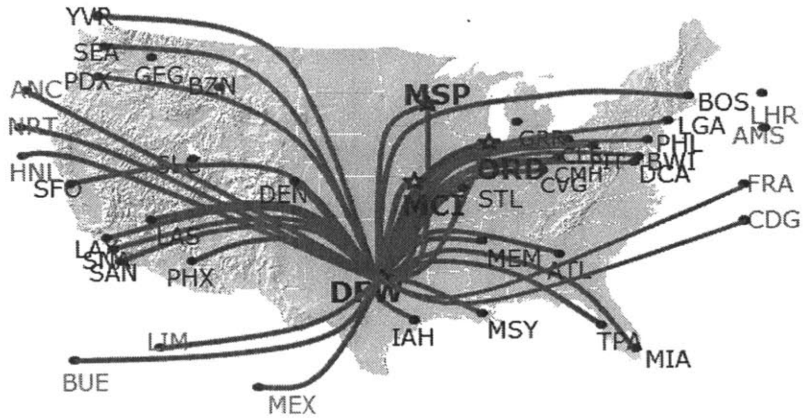

Figure 4.4: AL4 Network- DFW Hub

(Source: Belobaba, 2010)

The airlines offer three types of fare products: domestic restricted, domestic less-restricted and international restricted. A fare product consists of 10 fare classes (FC) with different restrictions applied to each fare class. The difference between the domestic fare products is the "aggressiveness" of the restrictions applied. Figures 5.1, 5.2 and 5.3 present the fare products offered in Network U10.

FC	AP	MIN3	CXL	SAT
1	0	0	0	0
2	0	0	1	0
3	3	0	1	0
4	7	0	1	0
5	7	1	0	0
6	7	1	1	0
7	14	1	1	0
8	14	1	1	0
9	14	1	1	1
10	21	1	1	1

Figure 5.1: Domestic Restricted

FC	AP	MIN3	CXL	SAT
1	0	0	0	0
2	0	0	1	0
3	3	0	1	0
4	7	0	1	0
5	7	0	1	0
6	7	0	1	0
7	14	0	1	0
8	14	0	1	0
9	14	0	1	0
10	21	0	1	0

Figure 5.2: Domestic Less-Restricted

FC	AP	SAT	CXL	MAX
1	0	0	0	0
2	0	0	1	0
3	3	0	1	0
4	0	1	0	0
5	3	0	1	0
6	7	1	1	0
7	14	1	1	0
8	14	1	1	1
9	21	1	1	1
10	21	1	1	1

Figure 5.3: International Restricted

(Source: Belobaba, 2010)

The restrictions included the fare products are advance purchase (AP), minimum three night stay (MIN3), cancellation fee (CXL), Saturday night stay (SAT) and maximum number of nights at destination (MAX). The cancellation penalty (fee) for domestic restricted, domestic less-restricted and international restricted is $200, $100 and $300, respectively. In line with current US airline industry practice, very few fare classes do not have a cancellation penalty associated with them, as in the real world only high fare class tickets are changeable without additional charge, or are fully refundable. As explained earlier, the cancellation penalty restrictions as well as all other restrictions described in Figures 5.1-5.3 are used by airlines for the purpose of preventing passengers with high WTP, mostly business, from buying down to lower fare classes.

3.4 Modelling Cancellation Behavior in PODS

Modeling passenger cancellation behavior can be done by allowing bookings to "leave" the system at some point in the time frame between the booking and departure times. Thus, each booking in PODS is assigned a cancellation probability at the beginning of each time frame that varies depending on the time frame. The findings in Illiescu et al. (2008) as well as industry experience suggest that passengers are more likely to cancel their tickets either shortly after they book or close to departure date, while in between they are less likely to cancel. This pattern resembles a "bath-tub" and hence the cancellation probabilities of each booking in PODS follow this pattern. An example of time frame dependent cancellation probabilities for a booking is demonstrated in Figure 6.1. In this example, a booking is made during TF8. In TF9 it is assigned a cancellation probability of 0.2, in between TF10 and TF14 it is assigned a significantly smaller cancellation probability of 0.01, and in TF15 and TF16 cancellation probabilities increase to 0.1.

TF	8	9	10	11	12	13	14	15	16
P(Cancel)	booked	0.2	0.01	0.01	0.01	0.01	0.01	0.1	0.1

Figure 6.1: Example of CXL Probabilities by Time Frame

In addition, each booking in PODS is assigned a cancellation probability depending on the passenger type, business or leisure, and whether or not a penalty (cancellation fee) is applied. The rationale here is that just as different types of passengers have different arrival or booking pattern depending on their characteristics and type of market, as explained earlier, they also have different cancellation patterns. Since business passengers' travel plans change more frequently compared to leisure passengers' travel plans, bookings of business passengers in PODS are assigned higher cancellation probabilities than bookings of leisure passengers. A booking that has a cancellation penalty associated with it is assigned a smaller cancellation probability than a booking without a penalty. The reasoning here is that a passenger will be less willing to change or cancel his ticket if there is a penalty involved. Figure 6.2 shows the probabilities used in PODS for simulation of a medium CXL rate scenario. In this scenario the total number of cancellations over all time frames and fare classes sums up to approximately 15% of all gross bookings.

Passenger Type	Penalty	TF after booking	Between	Last 2 TFs
Business	No	0.2	0.01	0.1
	Yes	0.1	0.005	0.05
Leisure	No	0.12	0.006	0.06
	Yes	0.06	0.003	0.03

Figure 6.2: CXL Probabilities by Passenger Type and Penalty

The table shows that a booking in a fare class without penalty is assigned a cancellation probability that is double the probability of a booking with penalty independent of passenger type. Also, leisure passengers with tickets that have no penalty associated have higher cancellation probabilities than business passengers with tickets that have penalty.

3.5 Passenger Cancellation Reporting in PODS

Every simulation run in PODS generates an output which includes a summary of the revenues, yield and load factor, among many other parameters, for each airline. PODS also provides a detailed report on forecasts, gross bookings and cancellations by class and time frame and the number of denied boardings by class (if there are any).

Figure 7.1 illustrates an example of gross bookings and cancelled bookings in each of the 16 time frames for one of the airlines in the simulation. Gross bookings are the number of bookings accepted in the reservation system across all fare classes. The figure shows that the number of cancelled bookings is more or less constant between TF2 and TF14, afterwards the number of cancelled bookings increases significantly due to the higher cancellation probabilities assigned to all bookings in the last two time frames before departure as explained earlier. The low number of gross bookings together with the higher number of cancelled booking in the last time frames means the reservation system actually loses more bookings than it gains new bookings and therefore the net bookings (i.e. gross bookings minus cancelled bookings) are negative and overall bookings-in-hand (BIH) decrease.

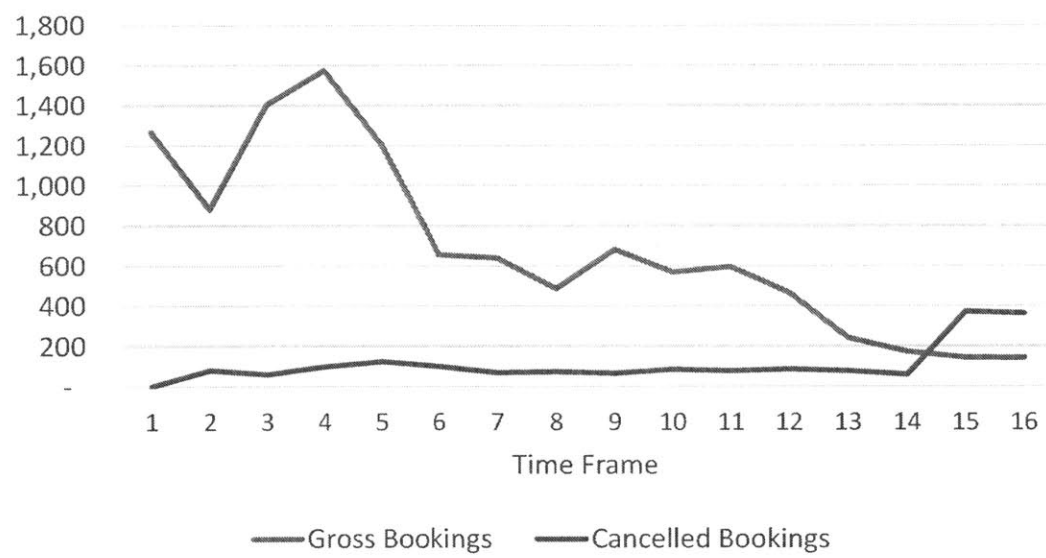

Figure 7.1: Example of Gross and Cancelled Bookings by Time Frame

Figure 7.2 presents an example of the distribution to fare classes of gross and net bookings for an airline in PODS.

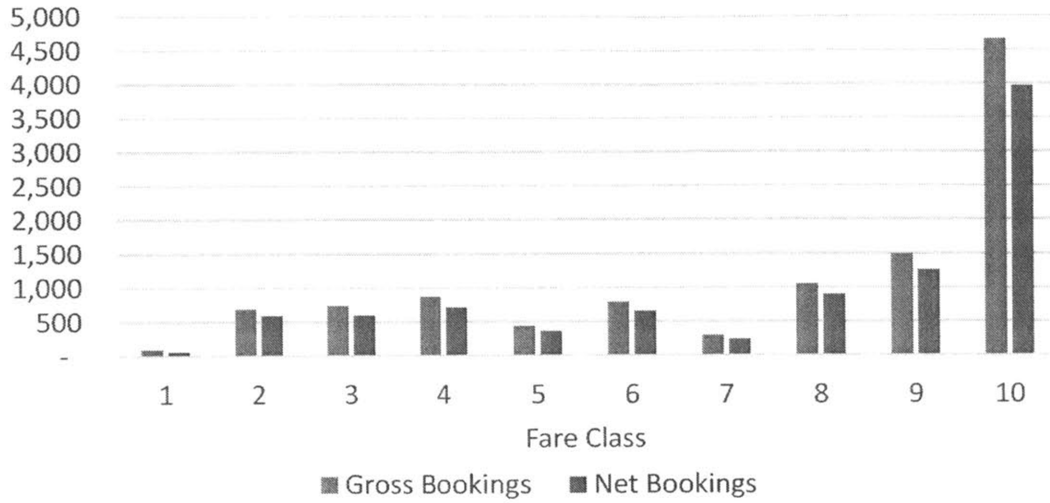

Figure 7.2: Gross and Net Bookings by Class

The figure shows that CL8, CL9 and CL10 bookings are about 65% of total gross bookings and that the cancellations are proportional to the number of gross bookings in each fare class. That is, the higher the number of bookings in a specific fare class, the higher the number of cancellations. Fare classes 2 through 10 have a cancellation rate of approximately 16%, while fare class 1 has a cancellation rate of 35%. FC1 does not have a cancellation penalty and is mostly booked by business passengers and hence has the higher cancellation rate.

3.6 Comparison between PODS and Industry Cancellation Rates

In order to validate that the cancellation probabilities assumed in PODS and the resulting cancellation rates more or less match cancellation rates in the real world, a dataset from a large North American (NA) airline containing the number of gross bookings by time frame as well as the number of cancellations by time frame and booked time frame was used. The data refers to single leg itineraries out of an airport in the US in the coach cabin for July 1st to July 28th 2014. Figure 8.1 summarizes the booking and cancellation data by market.

Market	GROSS BK	CXL	% CXL	NET BK
Domestic	221,689	30,466	13.74%	191,223
International	51,298	9,884	19.27%	41,414
Total	272,987	40,350	14.78%	232,637

Figure 8.1: North American Airline Cancellation Data

As in Network U10, the majority of the legs (approximately 81%) are domestic while the rest are international. The cancellation rate on the domestic legs is 13.74%, while on the international legs the cancellation rate is higher, 19.27%. Overall, the cancellation rate is about 14.78%. Figure 8.2 presents the gross bookings, cancellations and net bookings curves, cumulatively, based on the airline's data. It should be noted, that the airline uses 22 time frames in its booking horizon (versus 16 in Network U10). Note that the increase in the total number of cancellations is moderate up to time frame 19, while afterwards there is a jump in the number of cancellations. This trend is in line with the implementation of cancellation behavior in PODS which assumes higher cancellation probabilities for all bookings in the last two time frames as shown in Figures 6.1 and 6.2.

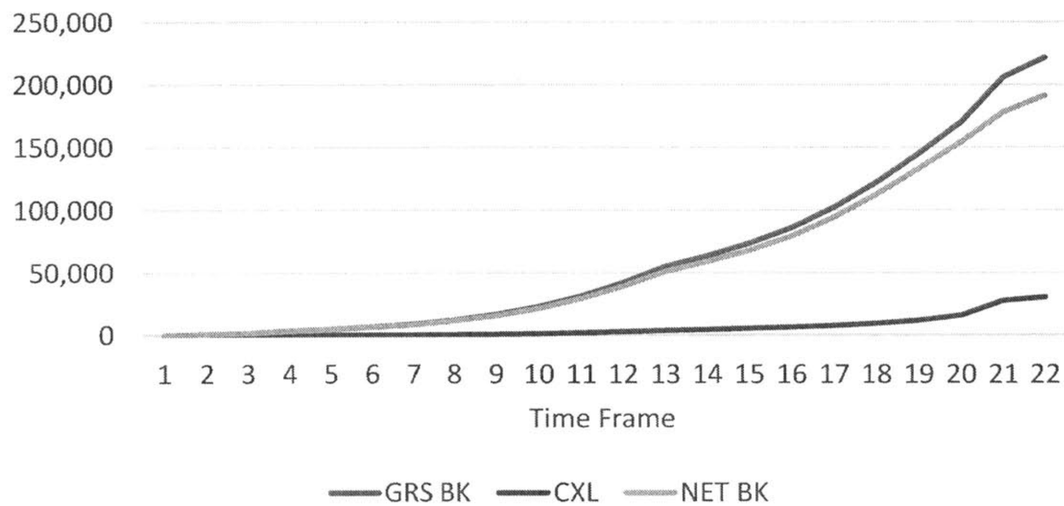

Figure 8.2: NA Airline Booking and Cancellation Curves

For the purpose of comparison between PODS' and actual airline cancellation rates, the NA airline's gross bookings from Figure 8.2 are now spread over 16 time frames (as in Network U10), and the of cancellations are calculated based on the cancellation probabilities shown in Figure 6.2. The gross bookings, cancellations and net bookings curves with PODS' cancellation probabilities are presented in Figure 8.3.

Overall, both the cancellation patterns and cancellation rates are similar. The cancellation rate is 13.16% and is close to the cancellation rate deriving from the airline data of 14.78% reported earlier. This conclusion serves as a validation that the cancellation probabilities used in this thesis can, in fact, be used to simulate "real world" passenger cancellation behavior and cancellation rates.

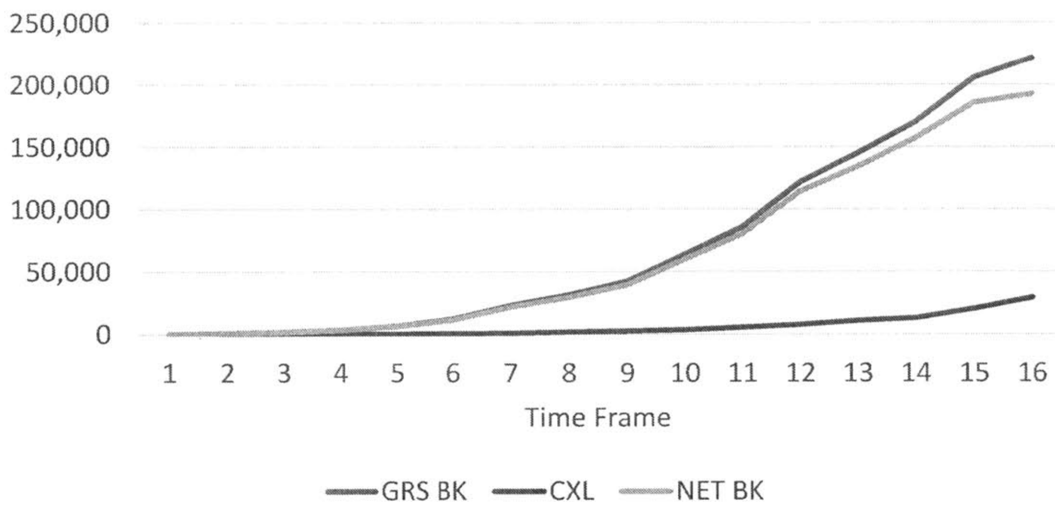

Figure 8.3: NA Airline Cancellations Based on PODS CXL Probabilities

3.7 Cancellation Forecasting Methods in PODS

In the overview of PODS' RM module earlier in this chapter, it was mentioned that the data the forecaster inputs into the RM seat allocation optimizer is based on the historical booking database that is constantly being updated as more bookings are accepted in the reservation system. However, this process ignored the cancellation phenomena and thus the impact of it on the accuracy of the forecaster. In order to allow the RM system in PODS to take into account the cancellation phenomena, the historical booking database was modulated to include both gross and net booking data at each time frame.

The cancellation forecasting methods in this thesis are an extension to the conventional historical based statistical models presented in Lemke et al. (2008) and Lawrence et al. (2003) where no-show numbers by booking class for future flights are forecasted by computing the mean no-show rate over a group of similar historical flights. However, the latter's model only looks at the total numbers of passengers booked and "no-show" passengers and disregards the time frame element in the computation of the no-show rates. In PODS, the cancellation estimates are forecasted by computing the mean number of cancellations at a path/class/time frame level. The overbooking method is based on the based on the deterministic overbooking model presented in Belobaba (2016) assuming the cancellations are known with certainty. Due to the current implementation of PODS, PNR based cancellation forecasting is not possible at this stage. In the sections below, several methods for cancellation forecasting will be proposed and discussed in detail. Each method has a different approach for cancellation estimates computation, forecast adjustment and compensation for cancellations. In addition, there will be two approaches for cancellation estimates aggregation and two approaches for adjustment of remaining capacity, or overbooking.

3.7.1 Cancellation Forecasting Method 1 (CM1)

CM1 is assumed to be the most basic method for cancellation forecasting in which the historical booking database records the number of net bookings at the end of each time frame and thus that is the only information used by the RM forecaster as an input to the seat allocation optimizer. Figure 9.1 provides an example of CM1's methodology. For the sake of simplicity, the example presents the historical booking data for a booking horizon of five time frames. Even though the number of gross bookings and cancelled bookings are presented here, the only information used for forecasting purposes is the "NET BK" data. At the beginning of TF1, the forecaster input into the optimizer equals 12 bookings, at the beginning of TF2 the input equals 8 bookings and so on. Unlike other cancellation forecasting methods described later in this thesis, there is no calculation of cancellation rates for either BIH or BTC and overbooking of remaining capacity is also not possible, as no data on cancelled bookings is recorded.

TF	GROSS BK	CXL IN TF	NET BK	NET BTC
1	4	0	4	12
2	3	1	2	8
3	3	1	2	6
4	3	1	2	4
5	3	1	2	2
TOTAL	16	4	12	

Figure 9.1: Example of CM1 Methodology

3.7.2 Cancellation Forecasting Method 2 (CM2)

CM2 is a more advanced method for cancellation forecasting that considers gross bookings and net bookings separately. The basic processes in CM2 follow the suggestion made by Chatterjee (2001) to separate between the BIH and BTC cancellations and therefore include:

- Forecasting gross BTC by path/class/time frame
- Forecasting cancellation of BTC by path/class/time frame
- Calculating cancellation of BIH by path/class/time frame

Cancellation estimation is complicated by the fact that some airlines only know the number of cancellations in a time frame, independent of the booked time frame, whereas others know both the number of cancellations by time frame and by time frame they were actually booked. In CM2 only the number of cancellations in a time frame is known and hence the cancellation estimate of BTC, *pcxltc*, by path/class in time frame *tf* (out of n time frames) is calculated at the beginning of each time frame as follows:

$$pcxltc_{p,c,tf} = \frac{\sum_{tf}^{n} bkgrs_{p,c} - \sum_{tf}^{n} bknet_{p,c}}{\sum_{tf}^{n} bkgrs_{p,c} + \sum_{tf=1}^{tf-1} bknet_{p,c}}$$

Where $bknet_{p,c,tf}$ is the total historical observations of net bookings for path *p*, class *c*, and $bkgrs_{p,c,tf}$ is the total historical observations of gross bookings for path *p*, class *c*. Alternatively, *pcxltc* can be translated to:

$$\frac{\text{Total number of cancellations left until departure (future CXL)}}{\text{Total gross bookings (Total GRS BK) - CXL already happened (past CXL)}}$$

In CM2, the cancellation estimate of BIH, *pcxlih*, is the same as *pcxltc.* Figure 9.2 provides an example for CM2's cancellation estimates methodology.

	START OF TF (FORECAST)				END OF TF		
TF	PAST CXL	TOTAL GROSS - PAST CXL	FUTURE CXL	BTC & BIH CXL RATE	GRS BK	CXL IN TF	NET BK
1	0	16	4	0.25	4	0	4
2	0	16	4	0.25	3	1	2
3	1	15	3	0.20	3	1	2
4	2	14	2	0.14	3	1	2
5	3	13	1	0.08	3	1	2
TOTAL					16	4	12

Figure 9.2: Example of CM2 Methodology

Following the calculation of the CXL estimates, BTC CXL estimates are used by the forecaster to scale down gross BTC forecast that is input into the optimizer to account for expected cancellation and BIH CXL estimates are used to reset, or overbook, the remaining booking capacity for each leg at the beginning of each time frame as follows:

- The product of BIH times the cancellation estimates for all path/classes using a leg is calculated. The result is the expected number of cancellations.
- An overbooking scaler (OBSCL) is chosen. The OBSCL could be any positive number. The higher the OBSCL, the more "aggressive" the overbooking will be. In this thesis, the OBSCL used are 0, 0.5, 1.0 and 1.5.

- The product of the expected cancellations and selected OBSCL is calculated and added to the remaining available capacity (i.e. physical capacity on leg minus BIH at the start of the time frame).

The following notation presents the calculation of booking capacity on leg l and time frame tf, $bcap_{l,tf}$:

$$bcap_{l,tf} = acap_l + obscl \cdot \sum_{p=1}^{npath} idpl_{p,l} \left[\sum_{f=1}^{nfcls} bkih_{p,c,tf} \left(pcxlih_{p,c,tf} \right) \right]$$

Where $npath$ and $nfcls$ represent the number of paths and fare classes, respectively, $acap_l$ is actual capacity on leg l, $idpl_{p,l}$ is a flag indicating whether (=1) or not (=0) path p uses leg l, $bkin_{p,c,tf}$ represents booking-in-hand in path p, class c, at the start of time tf.

Figure 9.3 shows a numerical example for the overbooking methodology of CM2 based on the example in Figure 9.2.

	START OF TF (FORECAST)					
TF	BIH	BTC & BIH CXL RATE	POTENTIAL OVERBOOKING= BIH* BIH CXL RATE *OBSCL			
			OBSCL=0	OBSCL=0.5	OBSCL=1.0	OBSCL=1.5
1	0	0.25	0	0	0	0
2	4	0.25	0	0.5	1	1.5
3	6	0.2	0	0.6	1.2	1.8
4	8	0.14	0	0.56	1.12	1.68
5	10	0.08	0	0.4	0.8	1.2

Figure 9.3: Example of CM2 Overbooking Methodology

3.7.3 Cancellation Forecasting Method 3 (CM3)

CM3, unlike CM2, assumes that cancellations are known by the airline both by time frame and by booked time frame. Thus, the cancellation estimate of BTC, *pcxltc*, by path/class in time frame *tf* is calculated as follows:

$$pcxltc_{p,c,tf} = 1 - \frac{\sum_{tf}^{n} bkad_{p,c}}{\sum_{tf}^{n} bkgrs_{p,c}}$$

Where $bkad_{p,c}$ is the total historical observations of bookings for path *p*, class *c* that survived until departure. Alternatively, *pcxltc* can be translated to:

$$\frac{\text{BTC that will be cancelled before departure}}{\text{gross BTC}}$$

The cancellation estimate of BIH, *pcxlih*, by path/class in time frame *tf* is calculated as follows:

$$pcxlih_{p,c,tf} = 1 - \frac{\sum_{tf=1}^{tf-1} bkad_{p,c}}{\sum_{tf=1}^{tf-1} bknet_{p,c}}$$

Or alternatively,

$$\frac{\text{BIH that have yet to cancel but will before departure}}{\text{BIH}}$$

Figure 9.4 presents an example for the calculation of BIH and BTC cancellation estimates under CM3.

	START OF TF (FORECAST)						END OF TF		
TF	GRS BTC	FUTURE CXL OF BTC	BTC CXL RATE	BIH	FUTURE CXL OF BIH	BIH CXL RATE	GRS BK	CXL IN TF	CXL OF BKD
1	16	4	0.25	0	0	--	4	0	2
2	12	2	0.17	4	2	0.50	3	1	1
3	9	1	0.11	6	2	0.33	3	1	1
4	6	0	0	8	2	0.25	3	1	0
5	3	0	0	10	1	0.10	3	1	0
TOTAL							16	4	4

Figure 9.4: Example of CM3 Methodology

Just as with CM2, BIH cancellation estimates are used to reset remaining capacity by adding the product of the expected cancellations times the OBSCL to the remaining available capacity and the BTC cancellation estimates are used to scale down gross BTC that are input into the optimizer. An example for CM3's overbooking methodology is also presented in Figure 9.3 as it is similar to the overbooking methodology of CM2.

3.7.4 Cancellation Forecasting Method 4 (CM4)

In CM4, both BIH and BTC cancellation estimates calculations are the same as in CM3, however the forecaster inputs into the optimizer gross BTC that are not scaled down to account for expected cancellations as in CM2 and CM3. As a result, the overbooking methodology of CM4 resets the remaining capacity based on the expected BIH and BTC cancellations.

The notation for the calculation of booking capacity on leg *l* and time frame *tf*, $bcap_{l,tf}$:

$$bcap_{l,tf} = acap_l + obscl \cdot \sum_{p=1}^{npath} idpl_{p,l} \left[\sum_{f=1}^{nfcls} bkih_{p,c,tf} \left(pcxlih_{p,c,tf} \right) + fcast_{p,c,tf} \left(pcxltc_{p,c,tf} \right) \right]$$

The only difference is the inclusion of *fcast* and *pcxltc* to the equation which represent the cancellation forecast of bookings-to-come in path *p*, class *c*, at the start of time *tf*.

Figure 9.5 shows an example of the overbooking methodology of CM4.

	START OF TF (FORECAST)							
TF	BIH	BIH CXL RATE	GRS BTC	BTC CXL RATE	POTENTIAL OB= (BIH* BIH CXL RATE + BTC * BTC CXL RATE) *OBSCL			
					OBSCL= 0	OBSCL= 0.5	OBSCL= 1.0	OBSCL= 1.5
1	0	--	16	0.25	0	0	0	0
2	4	0.50	12	0.17	0	2.02	4.04	6.06
3	6	0.33	9	0.11	0	1.49	2.97	4.46
4	8	0.25	6	0	0	1	2	3
5	10	0.10	3	0	0	0.5	1	1.5

Figure 9.5: Example of CM4 Overbooking Methodology

The different approaches of CM3 and CM4 with regard to the compensation for BTC cancellations and the forecaster input might result in different protection levels for the high fare classes and hence different overall results. Figure 9.6 presents an example of class protection levels for CM3 and CM4 while using different overbooking scalers for a given time frame. In the given example the BIH CXL RATE is 0.2, the BTC CXL RATE is 0.16 and the physical capacity of the aircraft is 100 and it is also assumed that the fare structure consists of only six fare classes (CL) and that the forecast is deterministic.

					CLASS PROTECTION							
	BIH	GRS BTC	FORECAST INPUT		OBSCL= 0		OBSCL= 0.5		OBSCL= 1.0		OBSCL= 1.5	
CM			3	4	3	4	3	4	3	4	3	4
			SCALE DOWN BTC	GRS BTC								
CL1	2	5	4	5	40	40	46	50	52	60	58	70
CL2	5	8	7	8	36	35	42	45	48	55	54	65
CL3	10	8	7	8	29	27	35	37	41	47	47	57
CL4	11	9	8	9	22	19	28	29	34	39	40	49
CL5	12	10	8	10	14	10	20	20	26	30	32	40
CL6	20	10	8	10	6	0	12	10	18	20	24	30
All	60	50	42	50								

Figure 9.6: Example of Class Protection with CM3 and CM4

As a reminder, CL1 protection level equals the remaining capacity on a flight leg as the RM system does not want to reject a demand for the highest fare class if there is any. One of the main differences between CM3 and CM4 is observed when OBSCL=0, or in other words, when no overbooking is applied. In this case, since no additional capacity is added either by CM3 or CM4 to the physical capacity, CL1 protection level is the same for both methods and equals 40. The difference in the protection levels (or the availability) of the lower fare classes increases the lower the class is. This is a direct result of the forecast input that CM3 uses which is a scaled down BTC versus the forecast input that CM4 uses which is gross BTC. Since the forecasts of CM4 are usually higher, the protections levels of the higher fare classes are higher meaning less availability left for the lower fare classes compared to CM3, again, if no overbooking is applied.

When overbooking is applied (i.e. OBSCL >0), CM4 adds more seats to the physical capacity than CM3 regardless of what the overbooking scaler is set to. Again, the overbooking mechanism of CM4 compensates for the expected cancellations of BIH and BTC and therefore will always add more capacity than CM3. The difference in the capacity adjustment of the two methods offsets the difference in forecast input and hence the difference in the availability of the lower fare classes is smaller than in the OBSCL=0 case.

CM4 becomes less aggressive than CM3 when OBSCL=1.5. In this case, the capacity added by CM4 versus CM3 exceeds the difference in the forecast input. The difference between the two methods' lower fare class protection can increase even further if OBSCL>1.5. At the same time, the number of denied boardings with CM4 will be higher than the number of denied boardings with CM3. Figure 9.7 summarizes the main concepts of all the cancellation forecasting methods discussed so far.

	CM1	CM2	CM3	CM4
Data used	**Net** bookings in each time frame	**Gross** bookings + CXL in each time frame	**Gross** bookings + CXL in each time frame + time frame CXL tickets were booked	**Gross** bookings + CXL in each time frame + time frame CXL tickets were booked
Cancellation estimates	No	**Same** CXL estimates for both BIH and BTC	**Different** CXL estimates for BIH and BTC	**Different** CXL estimates for BIH and BTC (same as CM3)
Overbooking remaining capacity	Not possible (No data on CXL)	**Product** of BIH CXL times the OBSCL	**Product** of BIH CXL times the OBSCL	**Product** of BIH & ***BTC*** CXL times the OBSCL
Forecaster input to optimizer	**Net** BTC	Scaled down BTC (after CXL)	Scaled down BTC (after CXL)	Gross BTC (before CXL)

Figure 9.7: PODS Cancellation Forecasting Methods Summary

3.7.5 Leg Based Cancellation Forecasting

In previous sections the cancellation forecasting methods proposed were based on calculating cancellation estimates on a path, fare class and time frame level. However, the capacity is based on a leg level. The difference in aggregation level of these two factors might lead to either too moderate or too excessive overbooking of flight legs as each leg in an airline's network serves multiple paths with different demand types (business or leisure) or volumes and at the same time different cancellation rates. It would be therefore appropriate to develop a cancellation forecasting method that is aggregated on leg, fare class, and time frame level. For this purpose, CM4 was extended to estimate cancellations on a leg level as well. *lcxlih*, the cancellation estimate of BIH by leg/class in time frame *tf* is calculated as follows:

$$lcxlih_{l,c,tf} = 1 - \frac{\sum_{n=1}^{npal_l} \sum_{i=1}^{tf-1} bkad_{idp_{l,n},c,i}}{\sum_{n=1}^{npal_l} \sum_{i=1}^{tf-1} bknet_{idp_{l,n},c,i}}$$

where $npal_l$ is the number of paths associated with leg *l* and $idp_{l,n}$ is the identity of the nth path using leg *l*. *lcxltc, t*he cancellation estimate of BTC by leg/class in time frame *tf* is calculated as follows:

$$lcxltc_{l,c,tf} = 1 - \frac{\sum_{n=1}^{npal_l} \sum_{i=tf}^{ntf} bkad_{idp_{l,n},c,i}}{\sum_{n=1}^{npal_l} \sum_{i=tf}^{ntf} bkgrs_{idp_{l,n},c,i}}$$

The booking capacity is then calculated as:

$$bcap_{l,tf} = acap_l + obscl \times \left[\sum_{c=1}^{nfcls} lbook_{l,c,tf} \left(lcxlih_{l,c,tf} \right) + \sum_{c=1}^{nfcls} \left(lcxltc_{l,c,tf} \right) \sum_{n=1}^{npal_l} fcast_{idp_{l,n},c,tf} \right]$$

where $lbook_{p,c,tf}$ represents the bookings on leg *l*, class *c* and time frame *tf*.

3.7.6 Apex Overbooking (APOB)

Apex overbooking is another approach for overbooking as an extension to the cancellation forecasting methods discussed in this thesis and it is based on a practice used by a North American airline. The idea behind APOB is to set the overbooking level (AU) to the minimum of that calculated by CM2, CM3 or CM4 and that calculated by APOB. APOB overbooking levels are generated by leg as follows:

- Obtain the average BIH by time frame using the accumulated booking history and cancellation history for the path and class using that leg
- Find the maximum (apex) BIH across all time frames
- For apex time frame and all earlier time frames, set APOB overbooking level (AU) to $CAP * \frac{MAX\ BIH}{NET\ BK\ AT\ DEPARTURE}$
- For each time frame after apex, set APOB overbooking level (AU) to $CAP * \frac{BIH\ TF(t)}{NET\ BK\ AT\ DEPARTURE}$

Where CAP represents the physical capacity of the aircraft.

Figure 9.8 presents the concept of APOB for an airline that uses CM2 or CM3 for a leg with a capacity of 100 seats and a booking horizon of eight time frames. Average net bookings on the flight are 100. The BIH curve represents the average BIH based on booking history and the apex is found at the start of time frame 6. The major assumption here is that there is only one apex over all time frames as, in reality, several apexes might be observed over the booking horizon. As expected, the AUs set by the OBSCL technique are low at the beginning of the booking horizon and increase as departure day approaches due to the linear correlation with the BIH. In addition, the higher the OBSCL, the higher the AU. On the other hand, the AU (=116) set by the apex technique is constant in the pre-apex time frames, while in the post-apex time frames the AU decreases from one time frame to the next one.

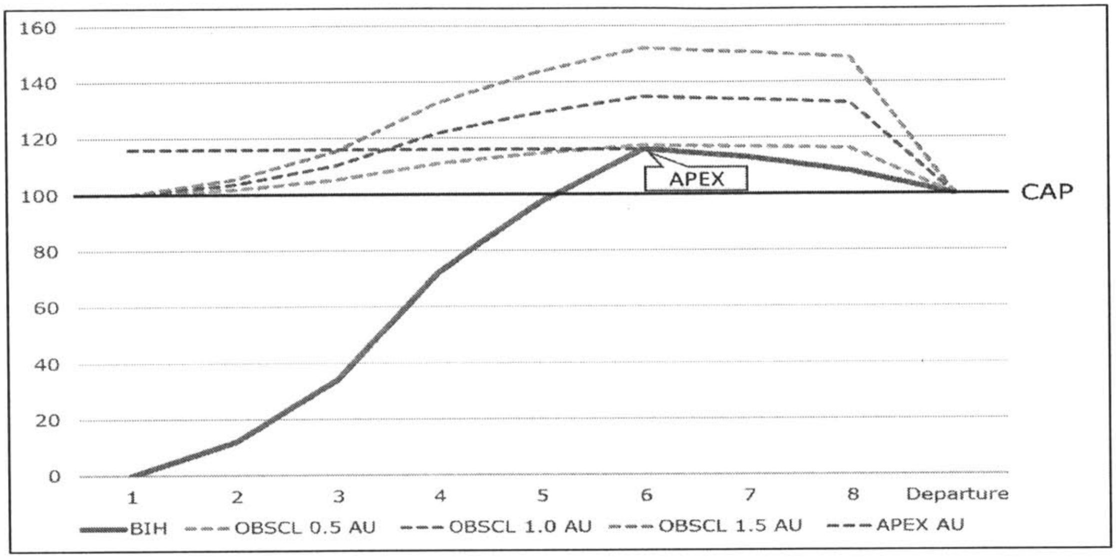

Figure 9.8: Example of OBSCL and APEX Overbooking Levels

Figure 9.9 illustrates the actual AU set by the APOB methodology with different OBSCL settings. The AUs presented in this figure are the minimum of the two AUs set by the OBSCL technique and the apex technique.

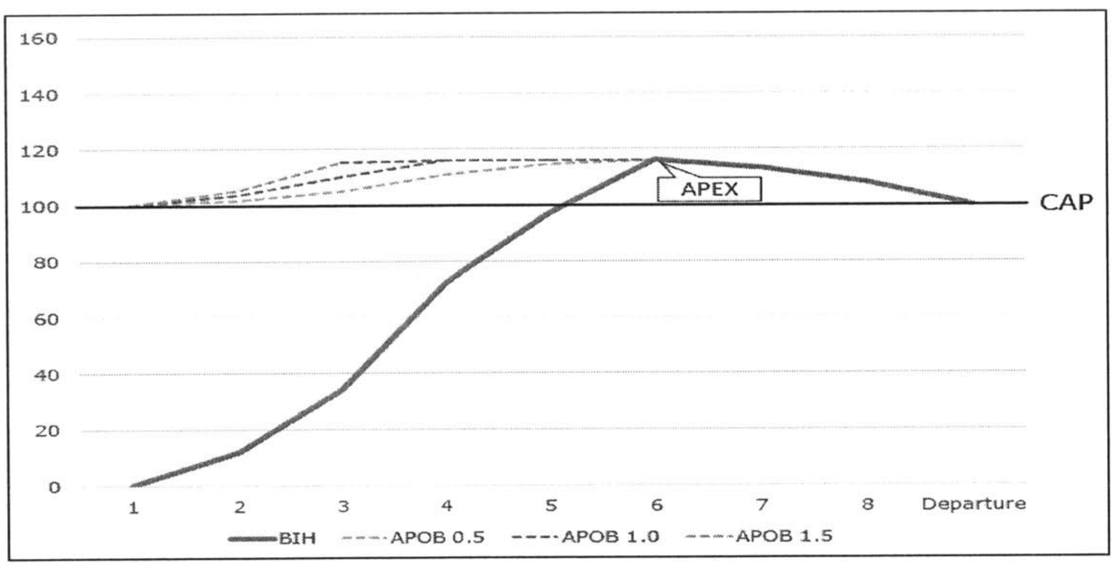

Figure 9.9: Example of Overbooking Levels with APOB

3.8 Summary

The research literature offers several approaches for the development of passenger cancellation forecasting and overbooking techniques. The methods in this thesis are based on the combination of these techniques together with practices from the US airline industry. The different methods included in this thesis represent only a fraction of the possible techniques for cancellation forecasting that are being used by airlines or other service based companies. However, measuring the value of each of these methods, or in other words, the revenue advantage (or lack thereof) of one method over the other is not a straightforward task for these companies. Using the PODS platform, the impact of each of these methods on the airlines' performance in terms of revenues, fare class bookings, load factors and yield will be analyzed. The results that will be presented in Chapter 4 are meant to provide a general perspective regarding the value of each of this methods rather than guaranteed results as the impacts of these methods also depend on myriad of factors, of which some the airlines do not have the ability to control. Therefore, the methods proposed will be tested in several different scenarios.

Chapter 4: PODS Simulation Results

This chapter presents and analyzes the results of a series of tests run in PODS. Each test consists of several scenarios which all have the same settings except for a few variables that will be modified in order to investigate their impacts on the RM system. In most tests, these variables will be changed for only one airline in the U10 network, namely AL1, while for all other airlines the settings will remain unchanged. The principal result measures that will be discussed are ticket revenues, number of denied boardings (DBs), net revenues, load factor and yield. The definition of each of these terms in the context of PODS is as follows:

- ***Ticket revenues***: All revenues collected from passenger bookings that survived to departure (net bookings). Revenues from cancelled bookings are not considered as ticket revenues.
- ***Denied boarding (DB)***: Passenger who did not fly (selected randomly) due to excess number of net bookings at departure over physical capacity.
- ***Net revenues***: Ticket revenues minus DB costs. DB costs are defined as the number of booked passengers who had been denied boarding times the input cost of a denied boarding (\$150/\$300/\$450)
- ***Load factor***: System load factor $\left(\frac{\text{Revenue Passenger Miles (RPM)}}{\text{Available Seat Miles (ASM)}}\right)$
- ***Yield***: Ticket revenues of passengers actually carried (net bookings minus DBs) divided by RPM.

Before presenting the results, the first sections in this chapter, 4.1.1 and 4.1.2, briefly describe the set of optimizers and forecasters used in the simulation tests.

4.1.1 RM Forecasters

The performance of the RM system is heavily reliant on the accuracy of the forecast that is fed into the optimizer. A highly sophisticated optimizer will not be able to produce optimal revenues if it is given a poor forecast. In this thesis, two types of forecasts are tested: standard path-class forecasting and hybrid forecasting (HF). Both forecasts are based on historical bookings of previous flights and current bookings of future flights.

Standard path-class forecasting is based on the "pick-up" forecasting methodology. Pick-up forecasting is a specific form of time-series forecasting that keeps track of previous unconstrained bookings and the changes in bookings over time, or in other words, the number of bookings that are "picked-up" from one time frame to the next. A BTC forecast is thus produced by adding the average picked-up bookings for the given time frame and the departure date from historical bookings. Standard path-class forecasting is an extension to the standard leg-class forecasting that is also implemented in PODS.

Hybrid forecasting is a more sophisticated forecasting technique that combines standard forecasting and "Q-forecasting", a technique developed by Belobaba and Hopperstad (2004) and mostly used in fully undifferentiated fare structures. Q forecasting is meant to prevent the spiral-down effect, as passengers prefer buying the lowest available fare when no restrictions are applied. Q forecasting incorporates passengers' WTP estimates to compute probabilities of sell-up from class Q (the lowest available fare class) to higher fare classes. Hybrid forecasting distinguishes between price-oriented demand, passengers who always purchase the lowest available fare, and product-oriented demand, passengers who are willing to buy higher fares due the reduced restrictions on them. The name "hybrid" derives from the fact that each demand group has a different booking behavior and hence different forecasts need to be utilized for each group. Pick up forecasting is used to forecast the product-oriented demand and Q-forecasting is used to forecast price-oriented demand. More details on Hybrid forecasting and Q forecasting can be found in Belobaba and Hopperstad (2004).

Another technique aimed at preventing the spiral down effect and commonly used with hybrid forecasting is the marginal revenue fare adjustment developed by Fiig et al. (2010). The technique addresses the concern that improving demand forecasts only by using WTP estimates is not enough to ensure revenue maximization especially for flights with capacity exceeding demand. RM optimizers need to also address the tendency of passengers to buy down in a given fare structure. The main assumption in fare adjustment is that input revenue values of lower classes should be reduced in order for the optimizer to reduce availability in those classes and close them earlier in the booking process and thus encourage sell-up, even when there are expected to be empty seats.

4.1.2 Seat Allocation Optimization

As a reminder, the optimizer in the RM system sets booking limits for each of the fare classes based on the input fares and forecasts. There are two general approaches for setting booking limits: leg-based controls and OD controls. The industry standard leg based optimizer called Expected Marginal Seat Revenue, or EMSR, was developed in Belobaba (1987b) and Belobaba (1989) as an extension to the seat inventory control models developed by Littlewood (1972), Buhr (1982), Richter (1982) and Wang (1982). The optimizer was refined to become EMSRb in Belobaba (1992). This approach assumes the demand for each fare class is independent of demand in other classes and that it can be described by a Gaussian distribution, based on historical data. Given the mean and the standard deviation for a given fare class, the expected marginal seat revenue is defined as the product of average fare of the fare class and the probability of the seat to be booked. The optimizer protects seats for a higher fare class only up to the point that the EMSR equals the fare of the next lower fare class. A full description of the EMSRb algorithm can be found in Belobaba and Weatherford (1996).

The relevance of leg based control has decreased in recent years due to the increasing use by large airlines of the "hub-and-spoke" type of operation that allows them to sell multi-leg itineraries. Thus, any seat on a flight leg can be shared by passengers from different fare classes that fly either on local or connecting itineraries. The OD (network) controls differ from the leg based controls by taking into account the network effects and hence they are set to maximize the revenues for the entire network and not for each individual leg separately.

In this thesis, three OD control mechanisms are simulated: Displacement Adjustment Virtual Nesting (DAVN), Probabilistic Bid-Price (ProBP), and Unbucketed Dynamic Programming (UDP). DAVN is a network OD control mechanism practiced by some of the largest airlines. It was developed by Wysong (1988) and Smith and Penn (1988). This OD control uses "value buckets" instead of fare classes for seat inventory management in which each OD itinerary/fare-type (ODF) combination is assigned to a revenue value bucket on each flight leg. The seat availability for a requested ODF combination is based on the availability of the relevant bucket on each leg of the passenger's itinerary. In addition, a penalty is applied to connecting fares to account for potential displacement of a local passenger on each flight leg. The network revenue of an ODF on the first leg of a connecting itinerary is equal to the itinerary fare minus the sum of the displacement costs of the other legs included in the itinerary. The displacement costs can be calculated either by using network optimization techniques or by simpler leg-based EMSR models. A more detailed description of DAVN can be found in Williamson (1992).

ProBP is an OD control mechanism that is based on the calculation of network bid prices which have the same role as the network displacement costs. The bid price is the sum of the marginal revenue values for an incremental seat on all legs of an itinerary and is the threshold number the airline uses to decide whether or not to accept a booking request. If the bid price is lower than the fare of the ODF requested, the request will be accepted. Bratu (1998) and Belobaba (2002) provide a complete description of ProBP.

The optimization models described so far are based on two main assumptions. First, the demand for each fare class is independent. Second, the arrival pattern of the demand is sequential. Dynamic programming (DP) is another approach used for calculating bid prices on a single leg developed by Lee et al. (1993) and Lautenbacher et al. (1999). Optimization models using DP differ from the static nature of traditional models by taking into account the interspersed arrivals of demand over the booking horizon and setting booking limits or bid prices within a time frame as a function of the

BIH and the remaining capacity. The major limitation of DP optimization is the size and complexity of the model formulation, which makes it infeasible to full network problems (due to computational constraints) and thus limits the use of DP to single leg problems in practice.

One of the network RM techniques using DP is called DAVN-DP. DAVN-DP employs the same steps as in standard DAVN expect for the last step in which EMSRb is replaced with a standard leg based DP to determine booking limits. Unbucketed DP (UDP) is a slightly modified version of DAVN-DP which relaxes the hard constraint of fixed (eight) virtual buckets used in DAVN-DP. The reader is referred to Diwan (2010) for further details on DP and especially DP based RM systems.

4.2 Test 1: Cancellation Forecasting Methods 1-3 (CM1-CM2-CM3)

The first test of this thesis is aimed at comparing the performance of the RM system using the first three cancellation forecasting methods developed in PODS, gaining insights whether or not one methodology is beneficial over the others and studying the effects of overbooking on the overall results. Thus in this test, AL1 will use CM1, CM2 and CM3 with different overbooking levels (OBSCL=0, 0.5, 1 & 1.5). AL1 using CM1 will be the base case scenario as CM1 does not actually forecast for cancellation. The settings for this test are as follows:

- Medium demand level (such that load factor is approximately 82%)
- Medium CXL Rate (with the same probabilities presented in Figure 6.1)
- Optimizer: AL1- DAVN, AL2- DAVN, AL3- EMSRb, AL4- DAVN
- Forecaster: All airlines use standard forecast
- AL2, AL3 and AL4 use CM2 with OBSCL=0.5 (moderate overbooking)

Figure 10 shows the total network ticket revenues for AL1 in each simulation. AL1 gains higher ticket revenues with CM2 and CM3 in all cases. When AL1 does not overbook (OBSCL=0), there are minor revenue gains of approximately 0.1% either with CM2 or CM3 in comparison with CM1. However, the higher the OBSCL is, the higher the ticket revenue gain CM2 or CM3 have over CM1. When AL1 overbooks moderately

(OBSCL=0.5), AL1 gains an additional 1.15% in ticket revenues with either CM2 or CM3. The ticket revenue gains over CM1 increase even more when AL1 overbooks more aggressively, in the magnitude of 2.08% and 2.07% when OBSCL=1.0, or 3.03% and 3.13% when OBSCL=1.5, for CM2 and CM3 respectively. The difference in revenue gains between CM2 and CM3 is marginal. Nonetheless, it should be reminded that ticket revenues are the revenues collected from all passenger bookings before taking into account the negative impact of DBs and their associated costs.

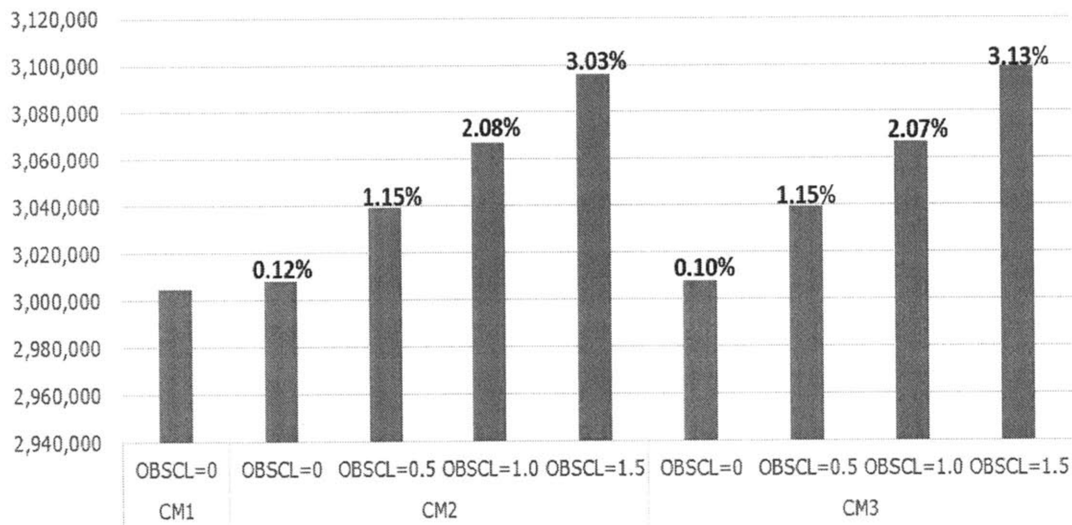

Figure 10: Test 1- Ticket Revenues & Percent Gain over CM1

Figure 11 shows the number of DBs for CM2 and CM3 under different OBSCL. CM1 is not included in this figure because overbooking cannot be applied with CM1 (and hence no denied boardings are possible). AL1 does have a higher number of DBs with CM2 and CM3 the higher the OBSCL is. However, while the number of DBs when OBSCL=0.5 is somewhat close to the US airline industry average number of eight DBs per 10,000 passengers booked[1], the number of DBs is nine times and even 20 times higher when OBSCL=1.0 or OBSCL=1.5, respectively. Evidently, using OBSCL that is greater than 1.0 is too aggressive and is far beyond what airlines in the real world would be willing to practice. In addition, the number of DBs with CM2 and CM3 is quite similar.

[1]http://www.rita.dot.gov/bts/sites/rita.dot.gov.bts/files/publications/national_transportation_statistics/html/table_01_64.html

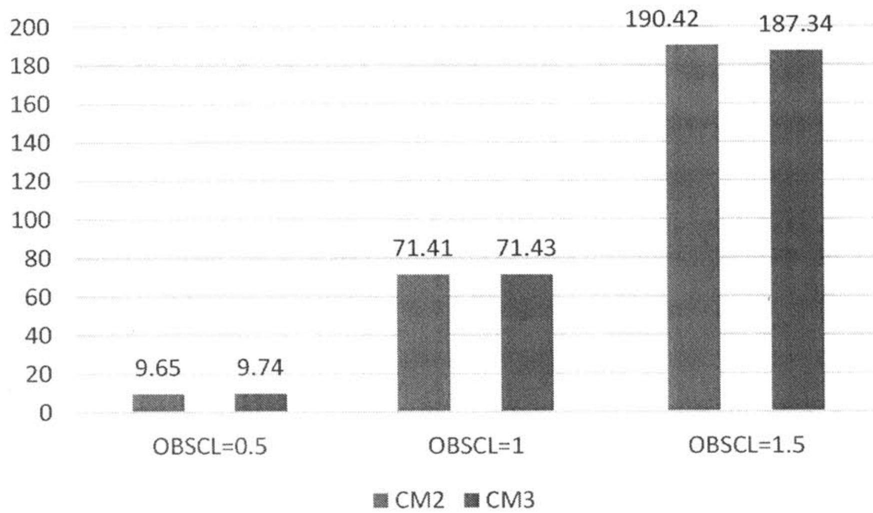

Figure 11: Test 1- Denied Boardings per 10K Passengers Booked

Figure 12 shows the net revenues of AL1. Again, net revenues are calculated as ticket revenues minus DB costs.

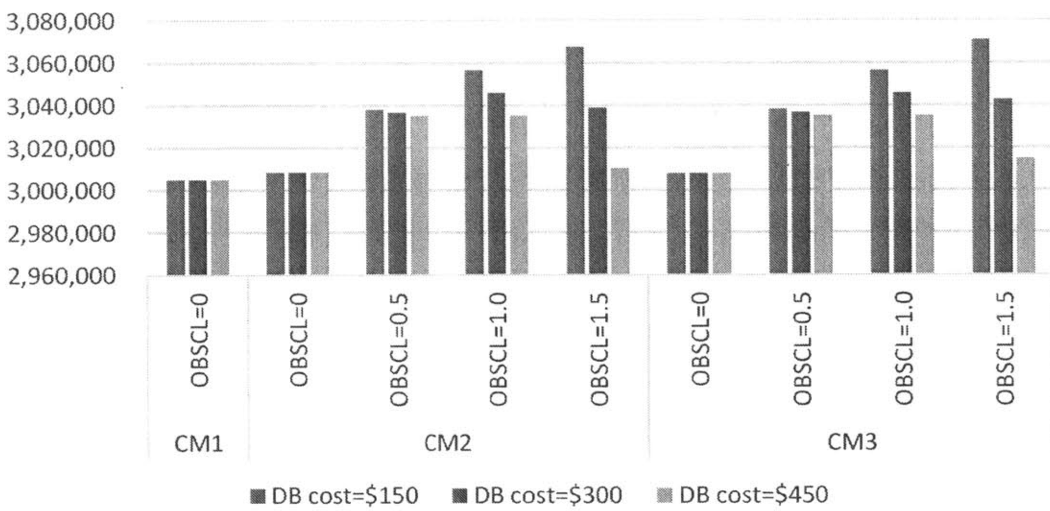

Figure 12: Test 1- Net Revenues

For CM1, CM2 and CM3 with OBSCL=0, the net revenues are equal to the ticket revenues in Figure 10 since there are no DBs. For the same reason, an increase in the DB cost has no impact on the net revenues. Just as with the ticket revenues, CM2 and CM3 have higher net revenues than CM1, no matter what OBSCL is used. Nonetheless, for each

OBSCL used the higher the input DB cost is the smaller the net revenue is, and the revenue gain over CM1 decreases as well. Figure 13 displays the percentage revenue gain of CM3 over CM1 in all scenarios. It should be noted that the revenue gains with CM2 are very similar and hence not presented.

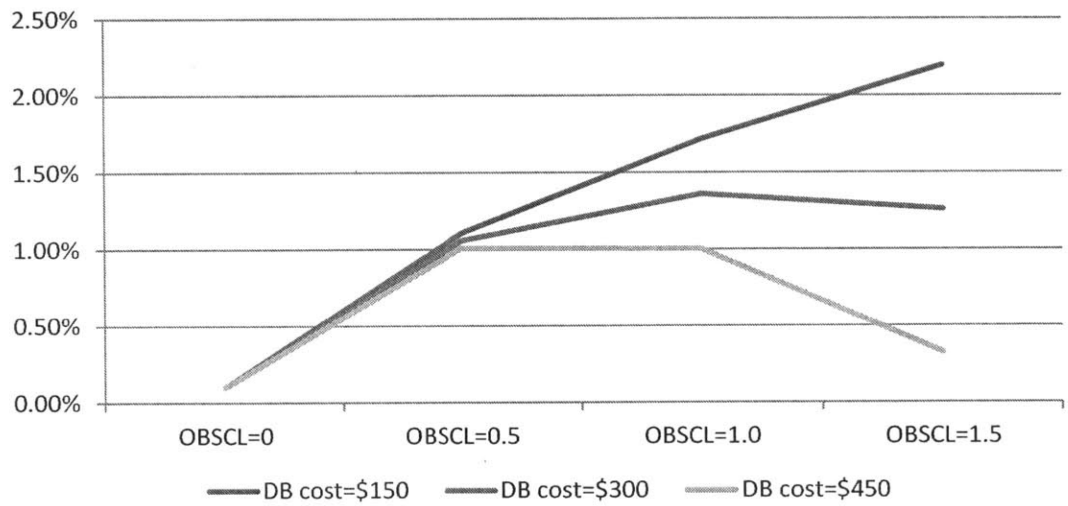

Figure 13: Test 1- CM3 Net Revenue Gain over CM1

When AL1 does not apply any kind of overbooking, the net revenue gains of CM3 over CM1 are equal to the ticket revenue gains of CM3 over CM1 presented in Figure 10. When OBSCL=0.5, the revenue gains over CM1 are approximately 1%, independent of the input DB costs. However, the differences in net revenue gains are much more apparent once AL1 chooses to use more aggressive overbooking measures. When OBSCL=1, CM3 net revenue gains over CM1 are 1.72%, 1.36%, 1.00% and when OBSCL=1.5, the revenue gains are 2.19%, 1.26% and 0.32% for DB costs of $150, $300 and $450, respectively. Based on these results, AL1 does benefit from overbooking in all cases. Nonetheless, the magnitude of the revenue gains from overbooking changes according to the DB costs and the aggressiveness of the overbooking. The higher these factors are, the lower the net revenue gains.

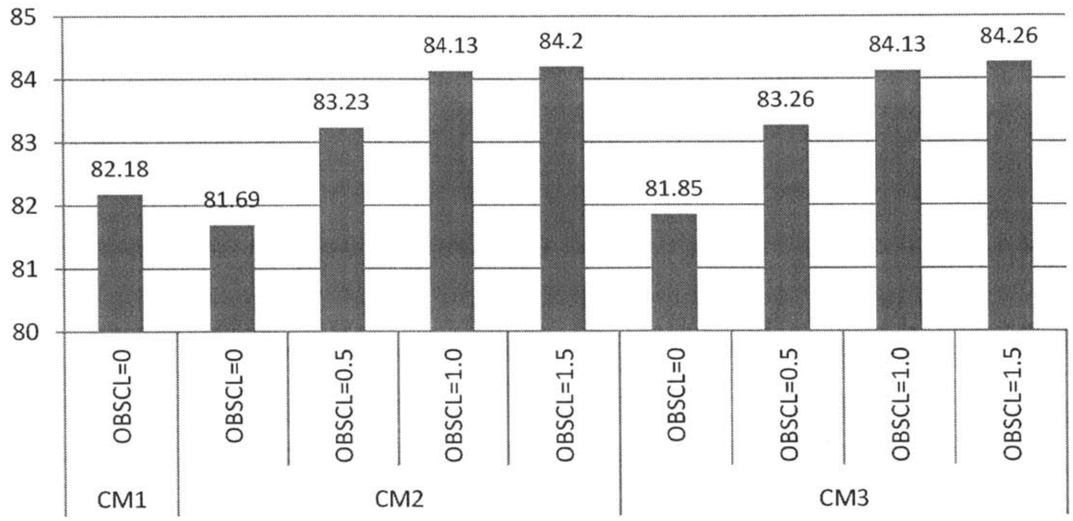

Figure 14: Test 1- Load Factors

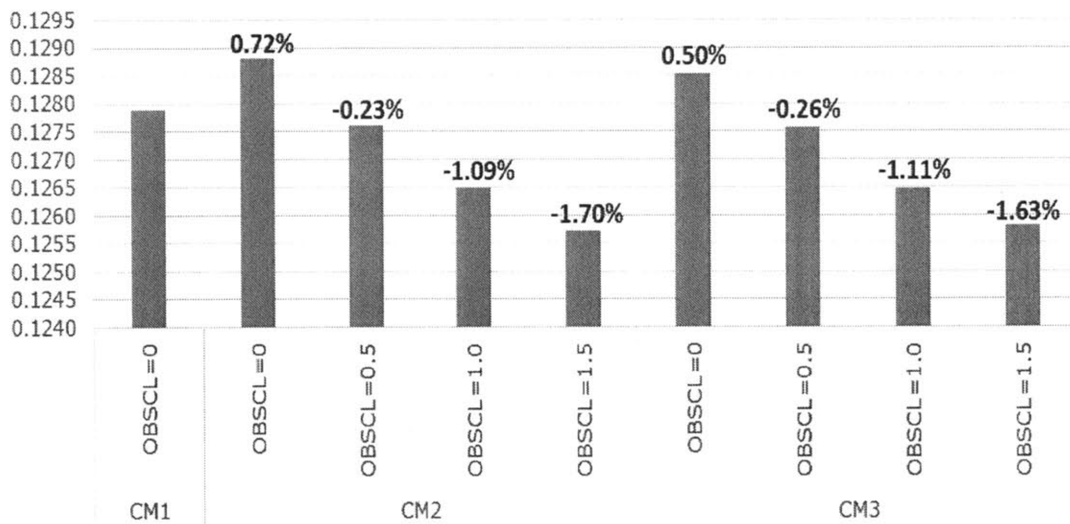

Figure 15: Test 1- Yields and Percent Change over CM1

AL1 load factors and yields in each scenario are presented in Figures 14 and 15. AL1 has the lowest load factors when it uses either CM2 or CM3 without overbooking. These load factors are lower CM1's load factor. As expected, the load factors increase as AL1 overbooks more aggressively. On the other hand, the yields decrease as load factors increase and hence the highest yields are observed when CM2 or CM3 are used without any overbooking. The results imply that when the RM system allows for overbooking, it accepts more passengers in low fare classes. When the overbooking is aggressive, the RM system accepts more low fare bookings and thus yields go down. The differences between the load factors and yields of CM2 or CM3 in each OBSCL level are marginal.

Figures 16 and 17 present the net bookings, or total number of bookings at time of departure, by class. Figure 16 shows the net bookings for CM1, CM2 and CM3 when no overbooking is applied. Figure 17 shows the net bookings for CM2 and CM3 when OBSCL=0.5 and OBSCL=1. Figure 18 shows the difference (in absolute terms) between the net bookings of CM2 and CM3 over CM1 with different OBSCL. When OBSCL=0, AL1 has less bookings in CL10 with CM2 and CM3 compared with CM1. The difference in net bookings in the higher fare classes is negligible. When OBSCL=0.5, AL1 accepts more bookings in CL6, CL9 and CL10 with CM2 or CM3 compared with CM1. The more aggressive the OBSCL is, more bookings are accepted in those classes. The difference in net bookings between CM2 and CM3 is marginal.

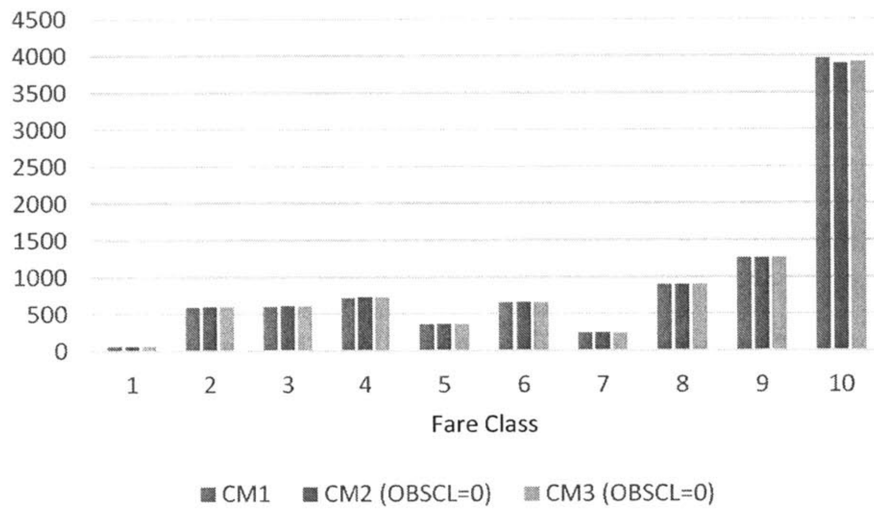

Figure 16: Test 1- Net Bookings of CM1/2/3 by Class- OBSCL=0

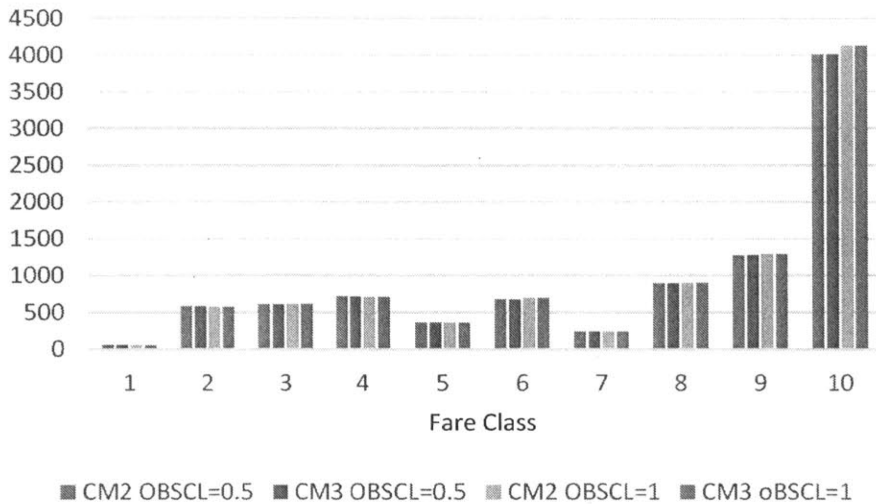

Figure 17: Test 1- Net Bookings of CM2/3 by Class- OBSCL>0

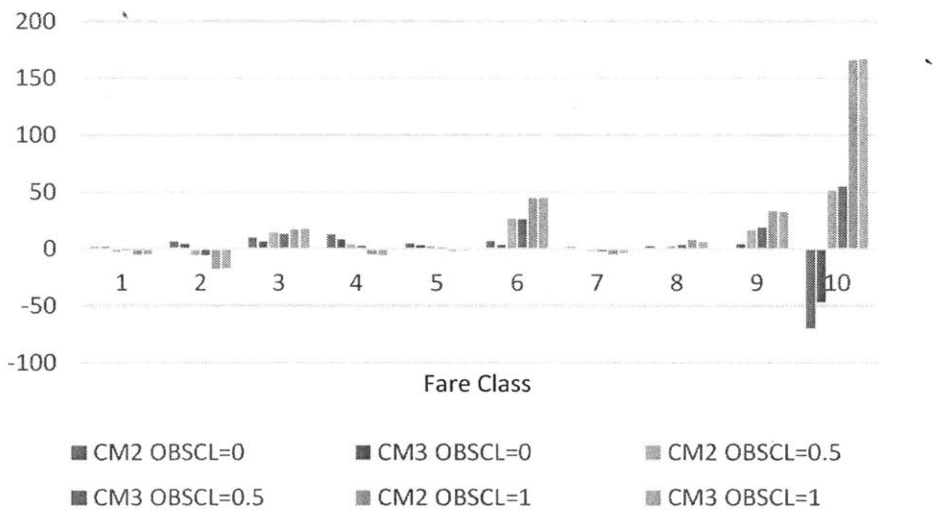

Figure 18: Test 1- Difference in Net Bookings by Class

As explained earlier in Chapter 3, the demand forecast is one of the main factors in determining the availability of fare classes for each future flight. Figure 19 displays the mean path forecast for CL10 bookings in the CM1, CM2 and CM3 with OBSCL=0 scenarios. It should be noted that CL10 has an advance purchase restriction and hence the forecast decreases from TF1 to TF9 and then the forecast is zero as CL10 is closed for bookings. In all the relevant time frames, the booking forecast for AL1 using CM1 is lower than the CM2 or CM3 forecasts, which are very similar. Studies on the accuracy of CM2 and CM3 have shown that both methods tend to (slightly) underestimate cancellation rates, or in other words the forecasted cancellation rates were lower than the actual rates, and hence the demand forecast fed into the optimizer was higher than what it should have been.

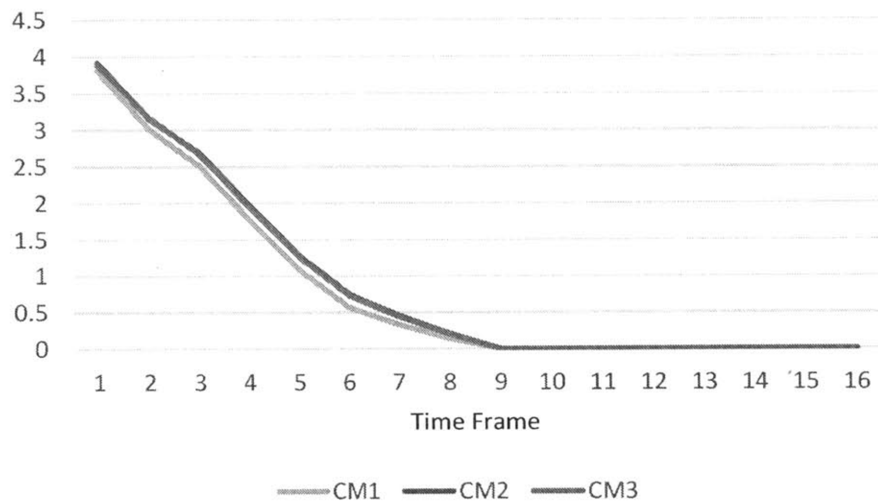

Figure 19: Test 1- CL10 Mean Path Forecast by Time Frame

Since AL1 is using the DAVN optimizer in this test, the forecasts are used to determine the EMSRc or "critical EMSR" calculated from the virtual buckets on each leg, which represents the minimum threshold for either accepting or rejecting booking requests. Figure 20 shows the EMSRc for AL1 using CM1, CM2 and CM3 with OBSCL=0 on each day in the booking horizon. The EMSRc of CM1 is lower than the EMSRc of CM2 or CM3 up to day 24 before departure and higher afterwards. The lower path forecasts (and EMSRc) allow the RM system to open more seats in lower fare classes when CM1 is employed. AL1 thus accepts more low fare class bookings and gains higher load factors and ticket revenues. The down side is that the yield is lower for CM1 as seen in Figure 15.

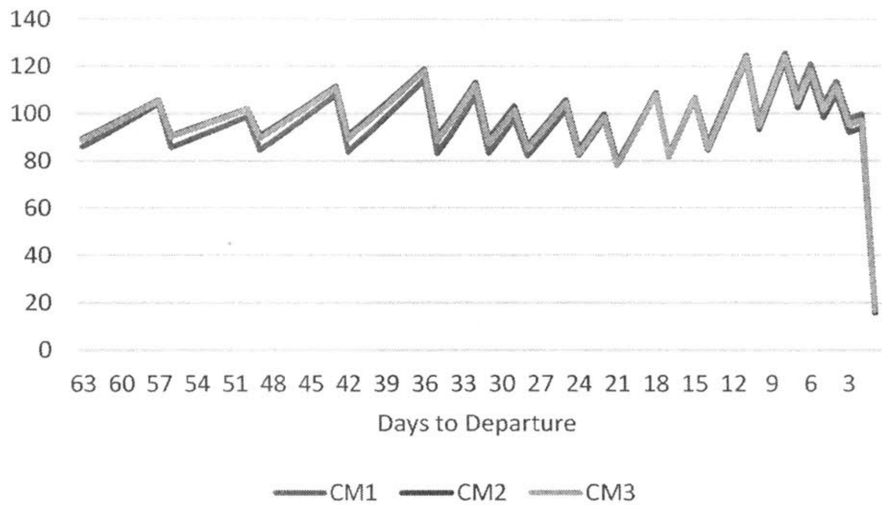

Figure 20: Test 1- EMSRc without Overbooking

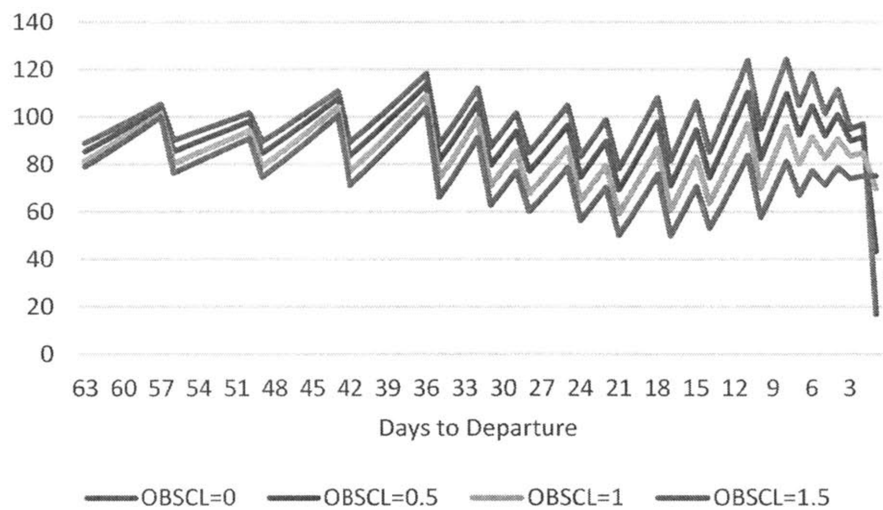

Figure 21: Test 1- EMSRc for CM3 with Different OBSCL

The interaction between overbooking and EMSRc is observed in Figure 21 which shows the EMSRc for AL1 using CM3 with different OBSCL. It should be noted that, the EMSRc of CM3 is very similar to EMSRc of CM2 (and thus not presented). Higher OBSCL is correlated with lower EMSRc and this explains the results presented in previous figures. The lower the EMSRc is early in booking horizon, the higher the availability is in the lower fare classes. The result is more bookings in low fare classes, higher load factors and higher overall ticket revenues, though yields are lower. Nonetheless, the number of denied boardings increases as well, with a negative effect on net revenues.

The relationship between the total number of bookings and EMSRc can be witnessed on the last day before departure in Figure 21. When AL1 does not overbook, it has the lowest number of bookings and due to the large number of cancellations right before departure, flights that were booked at physical capacity are now less full and the EMSRc drops sharply. The drop gets smaller the higher the OBSCL is. When AL1 uses aggressive overbooking parameters, the booking cancellations right before departure do not impact the EMSRc as much, as a large number of flights are still booked at or above physical capacity despite the cancellations.

The following points highlight the main conclusions arising for this test:

- Cancellation forecasting can help increase airline revenues (versus not forecasting cancellations as in CM1)
- The results AL1 achieves with CM2 are very similar to the results with CM3 despite the different methodologies, and thus neither method is clearly superior
- Overbooking can help increase the number of bookings and thus load factors and revenues
- Overbooking can have a negative effect on yields if as a result the RM system accepts more low fare class bookings
- Overbooking is beneficial in terms of net revenue if it is not overly aggressive and/or DB costs are low
- Aggressive overbooking may lead to large number of denied boardings which have negative effect on net revenues, especially if DB costs are high

4.3 Test 2: Cancellation Forecasting Methods 3 & 4 (CM3 & CM4)

The second test in this thesis compares the performance of the RM system when using the two more advanced cancellation forecasting methods described in Chapter 3, CM3 and CM4. These methods differ by the type of forecast that is being input into the optimizer and also the technique that is used for overbooking. Thus, in this test, AL1 will use CM3 and CM4 with different overbooking levels (OBSCL=0, 0.5, 1 & 1.5). AL1 using CM3 is the "base case" scenario for this test, as it has been discussed in detail in the previous test. The rest of the settings for this test are as follows:

- Medium demand level (such that load factor will be approximately 82%)
- Medium CXL Rate (with the same probabilities presented in Figure 6.1)
- Optimizer: AL1- DAVN, AL2- DAVN, AL3- EMSRb, AL4- DAVN
- Forecaster: All airlines use standard forecast
- AL2, AL3 and AL4 use CM2 with OBSCL=0.5 (moderate overbooking)

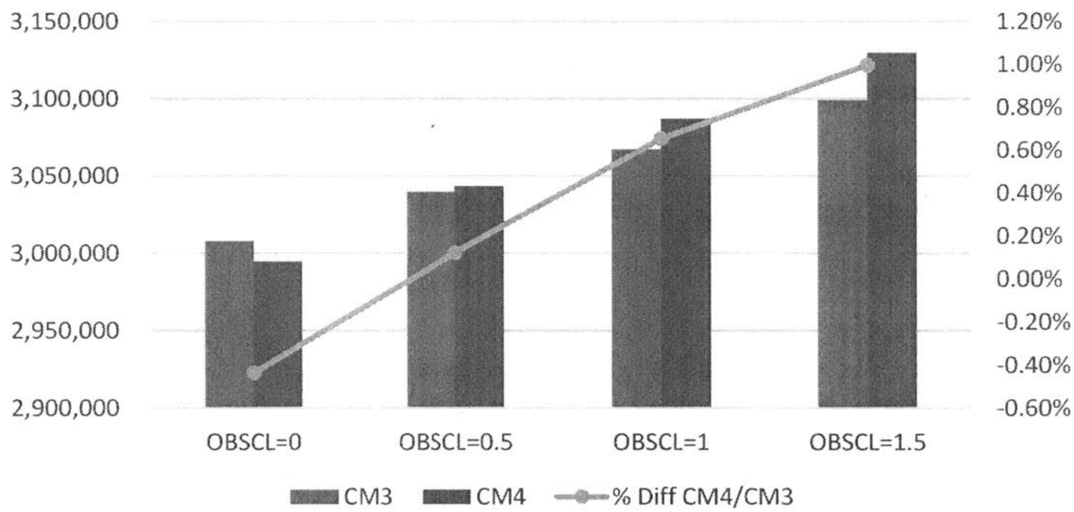

Figure 22: Test 2- Ticket Revenues & CM4 Revenue Gains over CM3

The ticket revenues of CM3 and CM4 are presented in Figure 22. When AL1 is not overbooking, CM3 has a revenue advantage over CM4. If OBSCL=0.5 or OBSCL=1.0, CM4 has a marginal revenue advantage of 0.12% and 0.16%, respectively, and a larger revenue gain of 1% if the overbooking is very aggressive (OBSCL=1.5).

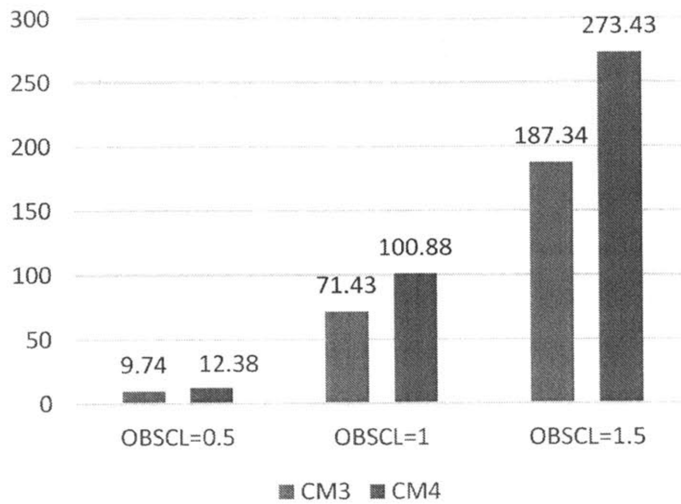

Figure 23: Test 2- Denied Boardings per 10K Passengers Booked

Figure 23 compares the number of DBs for AL1 using CM3 and CM4. In all scenarios, CM4 has a higher number of DBs and the difference, in absolute terms, increases the more aggressive the overbooking is. This is a direct result of the difference in the overbooking mechanisms of these methods. While CM3 only compensates for forecasted cancellations of BIH, CM4 compensates for forecasted cancellations of both BIH and BTC and thus AL1 usually sets higher AUs with CM4.

AL1's net revenues for CM3 and CM4 with overbooking are presented in Figure 24. Figure 25 shows the percentage revenue gain of CM4 over CM3 with different OBSCL. There are only two cases where CM3 has higher net revenues than CM4. The first case is when OBSCL=0. The second case is when OBSCL=1.5 and the DB cost input is $450. In this case the revenue gains of CM4 are completely diminished by the high number of DBs and the costs associated with them. In all other scenarios, the revenue gains of CM4 over CM3 range between 0.08% and 0.59%, depending, as usual, on the OBSCL and DB cost inputs.

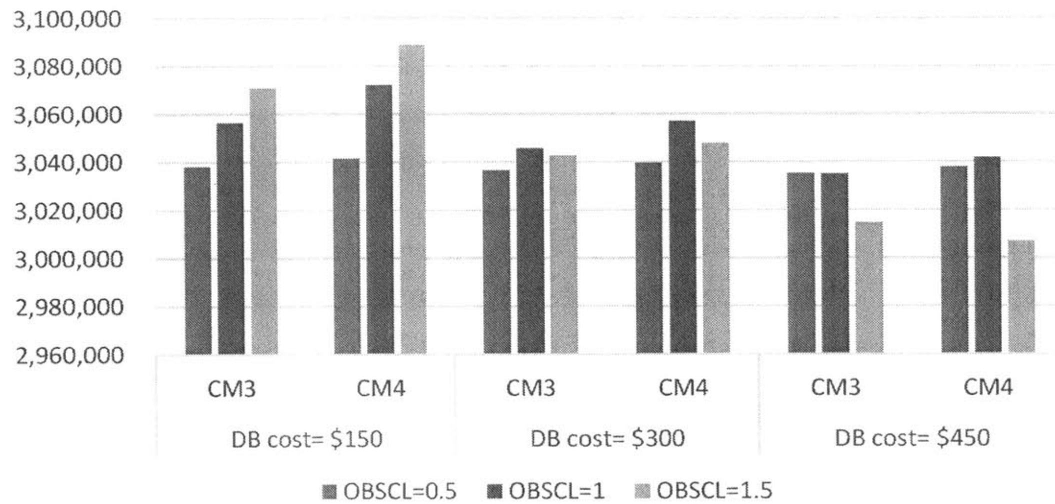

Figure 24: Test 2- Net Revenues

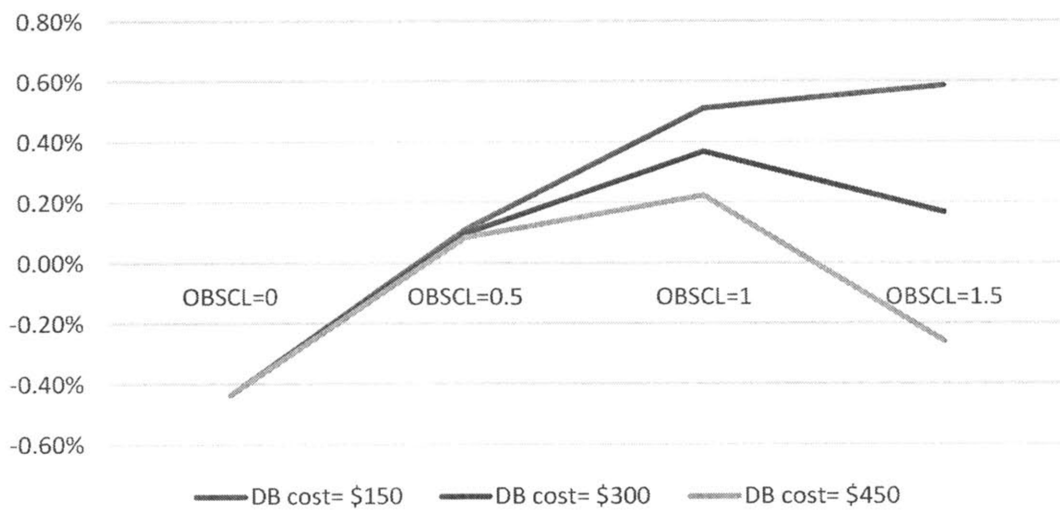

Figure 25: Test 2- CM4 Net Revenue Gains over CM3

The load factors and yields for AL1 using CM3 and CM4 are displayed in Figure 26 and Figure 27, respectively. AL1 has a load factor that is almost three percent points lower when using CM4 with OBSCL=0. This is a direct result of the difference in the forecasts that are fed into the optimizer as explained in Chapter 3 (Figure 9.6). However, the negative difference decreases up to the point where OBSCL=1, and then turns positive when OBSCL=1.5.

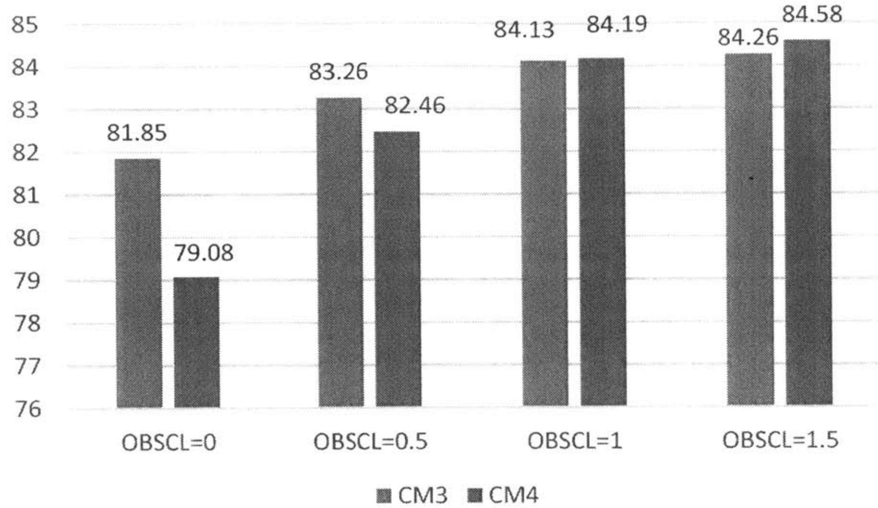

Figure 26: Test 2- Load Factors

The yields, in turn, have the opposite trend. CM4 has a significant 3.06% advantage in yield over CM3 when OBSCL=0 which turns into a 0.54% loss when the OBSCL=1.5. In addition, it is easy to notice that a change in OBSCL has a much bigger effect on the load factors and yields of CM4 than of CM3.

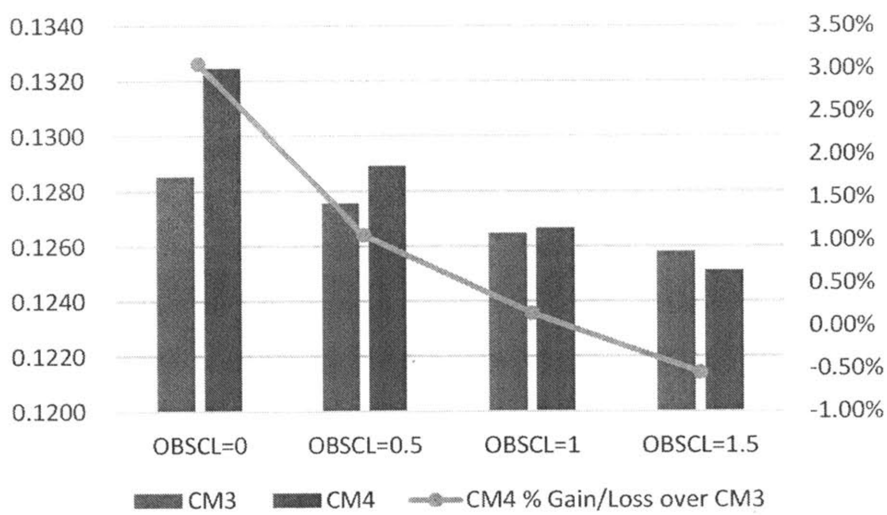

Figure 27: Test 2- Yields

The reasoning for the load factors and yields in this test can be found in Figure 28, where AL1's net bookings by class are shown. The most apparent differences are in the bookings of CL10. When OBSCL=0, AL1 has more bookings in CL10 with CM3 than with CM4 which brings CM3 to have higher load factors than CM4 but significantly lower yields. As OBSCL increases the difference in CL10 bookings decreases which, in turn, decreases the differences in load factor and yield. When OBSCL=1.5, CM4 has more CL10 bookings than CM3, which this time makes CM4 have a higher load factor but a lower yield.

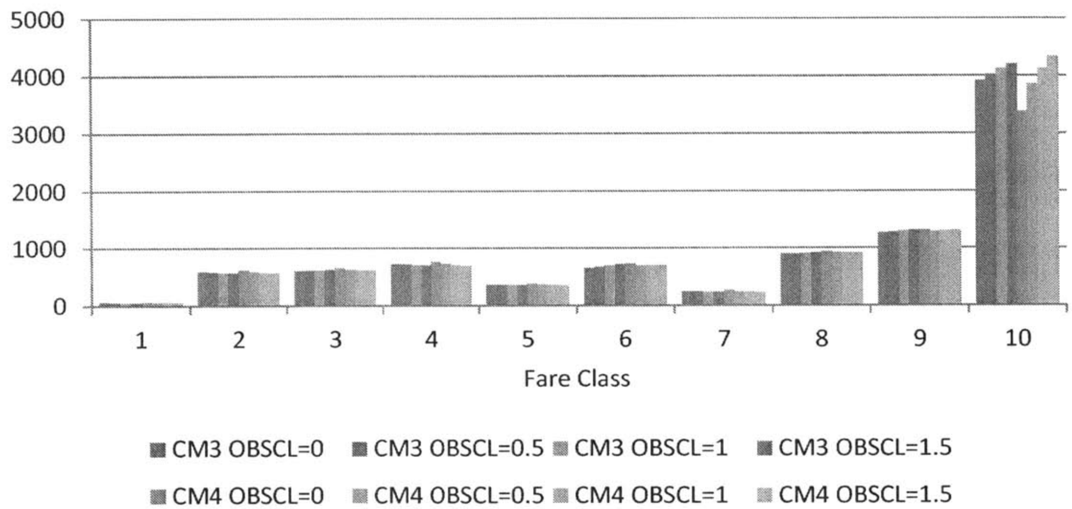

Figure 28: Test 2- Net Bookings by Class

The net bookings are an outcome of the EMSRc of CM3 (Figure 21) and CM4 (Figure 29). In general, OBSCL has a much bigger effect on EMSRc of CM4 than on that of CM3. When OBSCL=0, AL1 has higher EMSRc with CM4 than with CM3. This is a direct result of the difference in the forecast that is input into the optimizer. The optimizer uses a cancellation adjusted BTC forecast with CM3 and an unadjusted forecast with CM4. Higher bookings forecast means the optimizer allows fewer lower fare class bookings early in the booking horizon and hence more seats are available for late high fare class bookings.

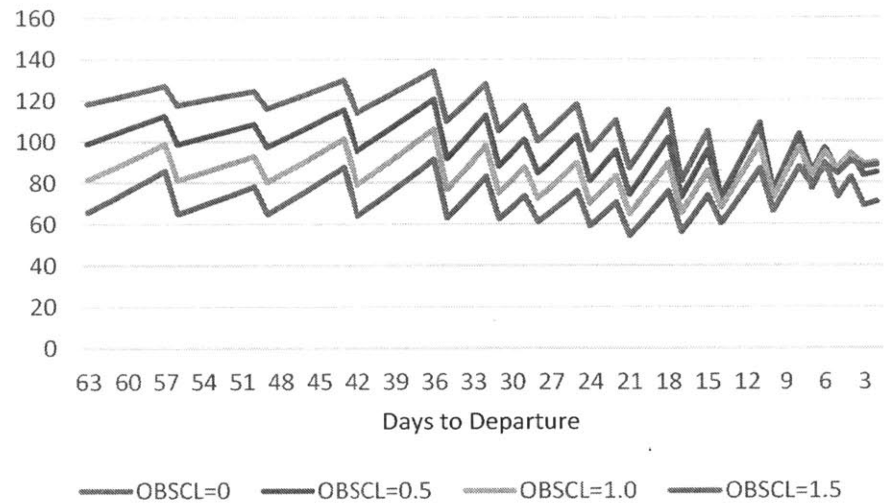

Figure 29: Test 2- EMSRc for CM4 with Different OBSCL

To conclude, these are the main findings from this test:

- CM4 ticket revenue gains are higher than CM3 when overbooking is applied
- CM4 has a significantly higher number of DBs than CM3 when OBSCL>0
- CM4 net revenues are higher than CM3 despite the higher number of DBs, except for two "extreme" scenarios where OBSCL=0 or OBSCL=1.5 and the DB costs are high ($450)
- There are large differences in number of bookings and thus load factors and yields due to the different approaches of these methods
- CM4 is more sensitive to changes in the OBSCL than CM3

4.4 Test 3: Competitive Environment

In the previous tests, the effects of AL1's RM cancellation and overbooking strategies on its own results were discussed in depth. The main assumption in those tests was that the other airlines in the network keep their RM strategies constant and do not respond with their own new strategies. Therefore, the third test in this thesis will address the implications of the RM strategy of the competitors on the results of AL1 while assuming AL1 keeps its RM strategy unchanged. The settings are as follows:

- Medium demand level (such that load factor will be approximately 82%)
- Medium CXL Rate (with the same probabilities presented in Figure 6.1)
- Optimizer: AL1- DAVN, AL2- DAVN, AL3- EMSRb, AL4- DAVN
- Forecaster: All airlines use standard forecast
- AL1 uses CM1 and CM3 with OBSCL=0, 0.5, 1.0

The scenarios in this test are divided into two types: the competitive scenario and the non-competitive scenario. The competitive scenario is the framework that was used in previous tests where AL2, AL3 and AL4 used CM2 and OBSCL=0.5. In the non-competitive scenario AL2, AL3 and AL4 use CM1 and thus do not forecast for cancellations and do not have the ability to overbook accordingly.

AL1 ticket revenues are presented in Figure 30. In all scenarios, AL1 gains lower revenues when the competitors are using a competitive cancellation strategy. AL1's ticket revenues decrease by approximately -0.43% when using CM1 and between -0.58% and -0.79% when using CM3 and different OBSCL. The DBs are presented in Figure 31. The difference in the number of DBs is marginal when OBSCL=0.5, nonetheless it increases when the OBSCL is higher than 0.5.

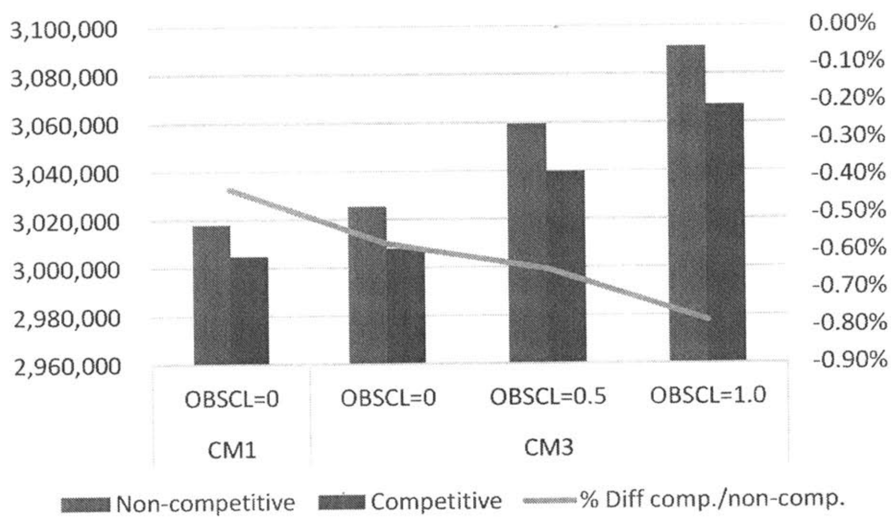

Figure 30: Test 3- AL1 Ticket Revenues

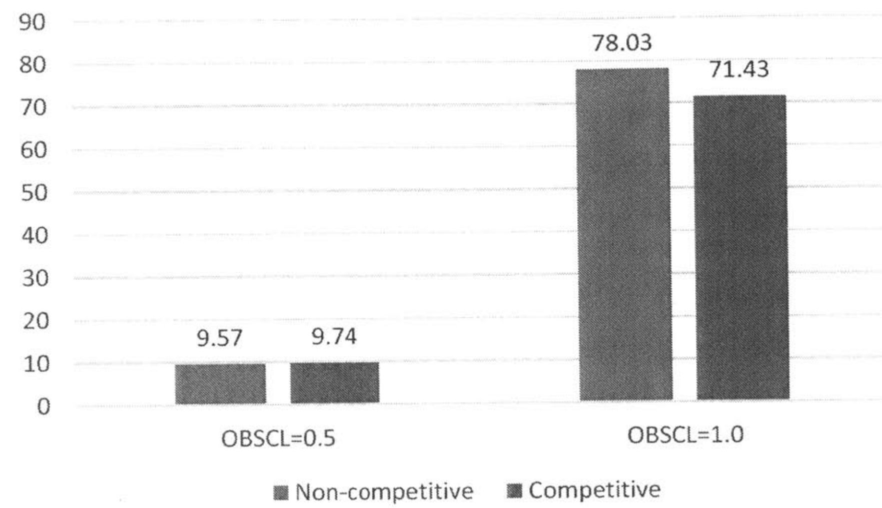

Figure 31: Test 3- AL1 Denied Boardings per 10K Passengers Booked

Taking into account the DB costs results in the net revenues shown in Figure 32. AL1 has lower net revenues in the competitive scenario in both OBSCL scenarios. This is a direct result of the lower ticket revenues AL1 gains in the competitive scenario as seen in Figure 30. Figure 33 compares the advantages of cancellation forecasting and overbooking under both scenarios. The net revenue gains over CM1 are approximately 0.2 and 0.3 percent points higher in the non-competitive case (compared with the competitive case) when OBSCL=0.5 and OBSCL=1, respectively.

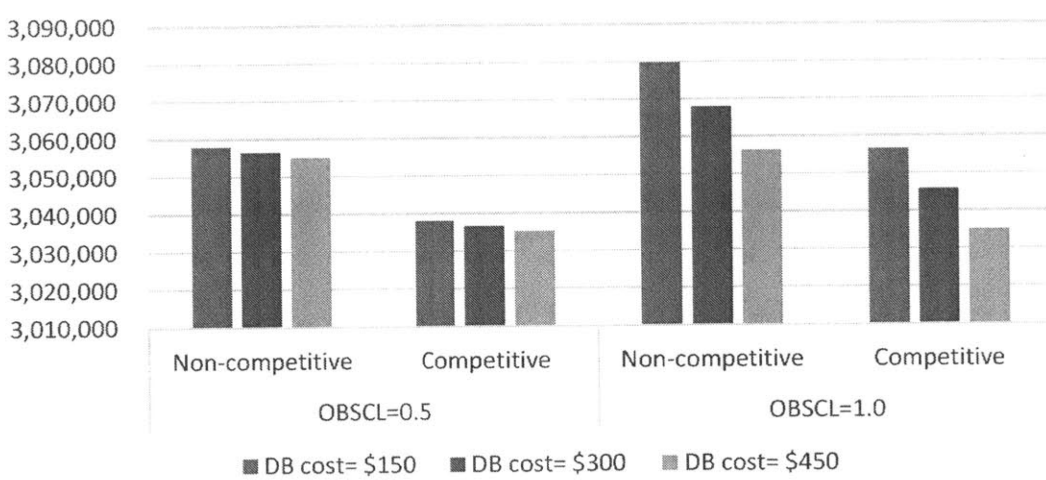

Figure 32: Test 3- AL1 Net Revenues

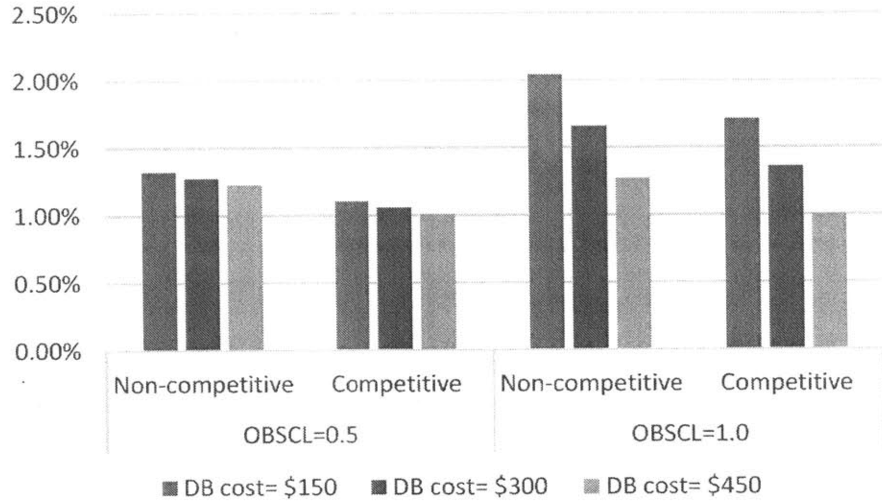

Figure 33: Test 3- % Net Revenue Gain over CM1

The load factors and Yields for AL1 are presented in Figure 34 and Figure 35, respectively. AL1 load factors are lower in the competitive scenarios. The interesting phenomenon in this test is that yields are also between 0.1% and 0.3% lower, depending on the OBSCL used.

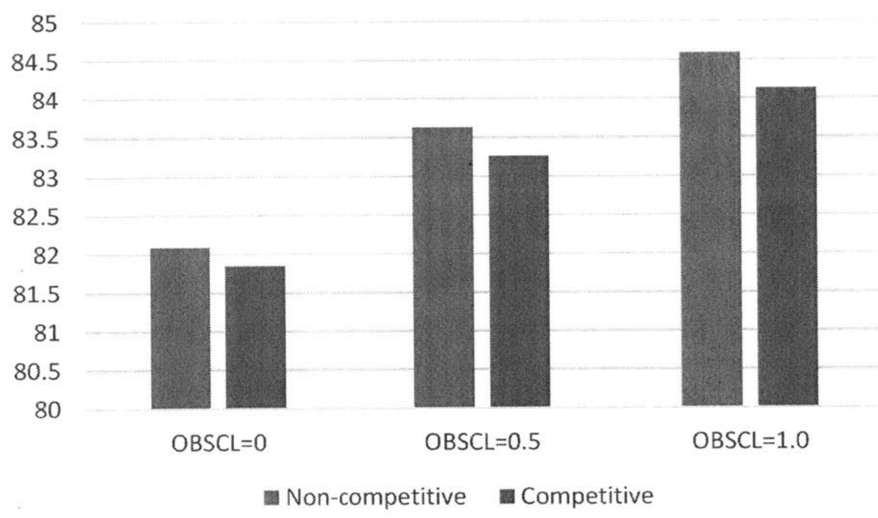

Figure 34: Test 3- Load Factors

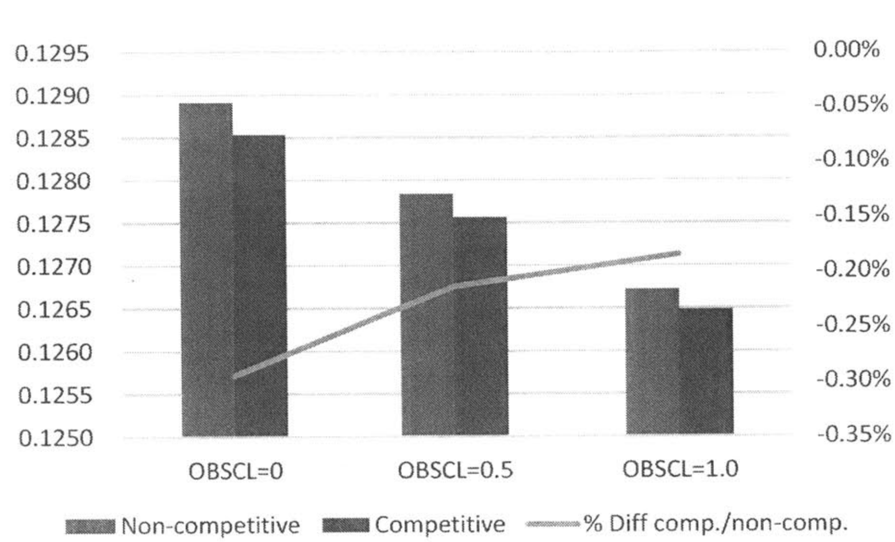

Figure 35: Test 3- Yields

Figure 36 displays the difference (competitive minus non-competitive) in net bookings between the two scenarios. The figure shows that no matter what the OBSCL is set to, AL1 has more CL10 bookings in the competitive scenario and, in turn, less bookings in all higher fare classes. The higher the OBSCL is the smaller the difference in CL10 bookings as well. In addition, AL1 has overall less net bookings in the competitive sub-scenario. Based on this finding, it is possible to explain the trends in Figures 34 and 35. In the competitive sub-scenario, all other airlines use a more advanced cancellation forecasting method and also overbook moderately which means they now have more seats available especially in lower fare classes. Consequently, passengers who would otherwise book with AL1 are now "spilled" to the competition. In order to reduce the spill, AL1's RM system is being less restrictive and opens more seats in low fare classes to encourage high WTP passengers to buy-down and book in lower fare classes. This is also reflected in the EMSRc of AL1 presented in Figure 37[2]. The competitive scenario has lower EMSRc than the non-competitive scenario for all OBSCL scenarios.

[2] For sake of simplicity, only the EMSRc when OBSCL=0 or OBSCL=0.5 is presented as the trends are similar with higher OBSCL

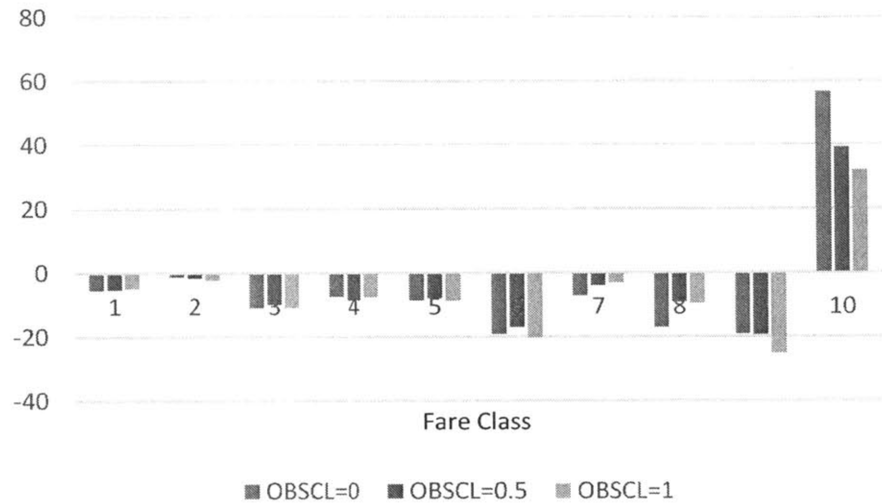

Figure 36: Test 3- Difference in Net Bookings by Class

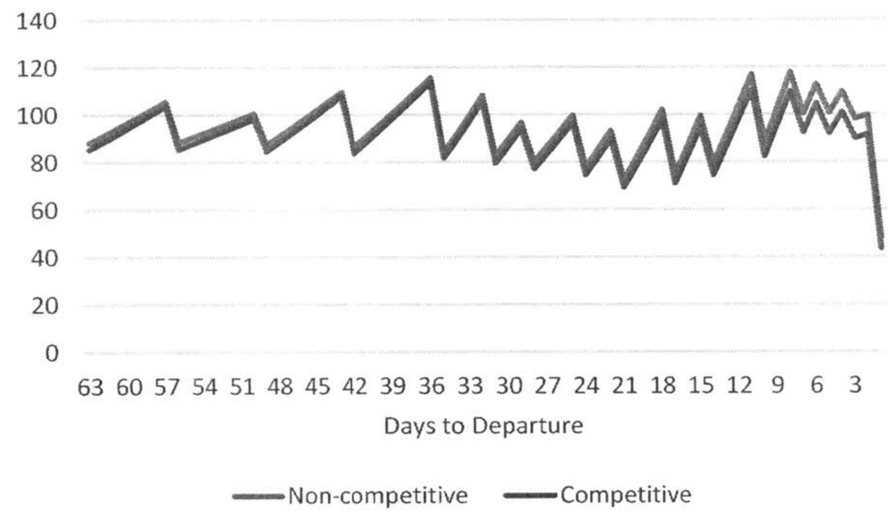

Figure 37: Test 3- EMSRc

In conclusion, this test compared the results of AL1 under two different scenarios of RM cancellation sophistication used by the competition. As expected, when the competitors use more sophisticated cancellation forecasting and overbooking strategies, AL1's ticket revenues and bottom line (net) revenues decrease. When the competitors open more seats, especially in low fare classes, AL1 lost high yield passengers to the competitors. At the same time, AL1's RM system responded by opening more seats in

CL10 which led to a spiral down effect. Nonetheless, AL1 can increase its net revenues by using more advanced cancellation forecasting and overbooking strategies (rather than not forecasting cancellations or overbooking at all) and the expected net revenue gains could range between 1% and 1.7%.

4.5 Test 4: Aggregation of Cancellation Estimates

The fourth test in this thesis explores the differences in the performance of the RM system using two different approaches for generating cancellation estimates. The first approach aggregates cancellation estimates on a path/class level, just as in all tests so far, and the second approach aggregates cancellation estimates on a leg/class level, as presented in section 3.7.5. The rest of the settings for this test are:

- Medium demand level (such that load factor will be approximately 82%)
- Medium CXL Rate (with the same probabilities presented in Figure 6.1)
- Optimizer: AL1- DAVN, AL2- DAVN, AL3- EMSRb, AL4- DAVN
- Forecaster: All airlines use standard forecast
- AL1 uses CM4 and OBSCL=0, 0.5, 1.0
- Other airlines use CM2 and OBSCL=0.5

AL1's ticket revenues are presented in Figure 38. As expected, there is no difference in ticket revenues when OBSCL=0 as the cancellation estimates are not used for either for adjusting the BTC forecast or for overbooking. When OBSCL>0, AL1 has higher ticket revenue gains when AL1 uses cancellation estimates aggregated on a leg/class level rather than a path/class level. The leg estimates' revenue gains are between 0.17% and 0.25% higher, depending on the OBSCL. The difference between the two aggregation approaches is large when looking at the DB numbers in Figure 39. The leg estimates approach has more DBs than the other approach with both OBSCL used. The bigger the OBSCL, the bigger the difference is in absolute terms.

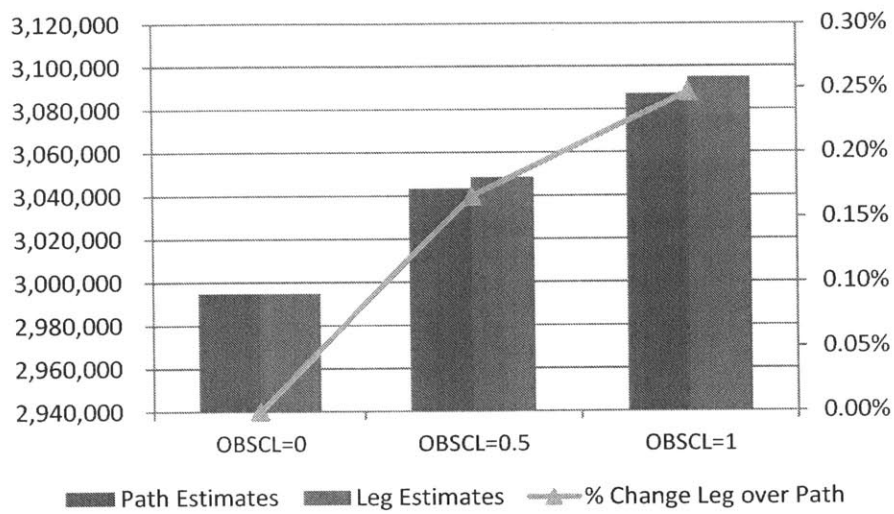

Figure 38: Test 4- Ticket Revenues

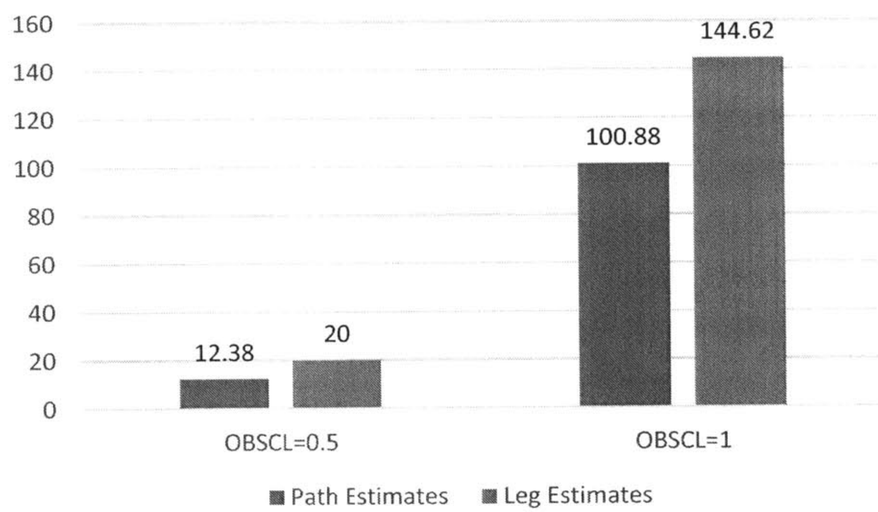

Figure 39: Test 4- Denied Boardings per 10K Passengers Booked

Figure 40 shows the net revenues of AL1 with both aggregation approaches and Figure 41 shows the percent revenue gain of the leg estimates over the path estimates. AL1 gains higher net revenues when using leg estimates and OBSCL=0, independent of the DB costs, or when OBSCL=1 and the DB costs are low ($150). In these cases the revenue gains range between 0.03% and 0.13%. When OBSCL=1 and the DB costs are high (>$300), the leg estimates loss up 0.4% in revenues compared with the path estimates. This is a direct result of the significantly higher number of DBs with the leg estimates approach and the reasoning will be presented later in this section.

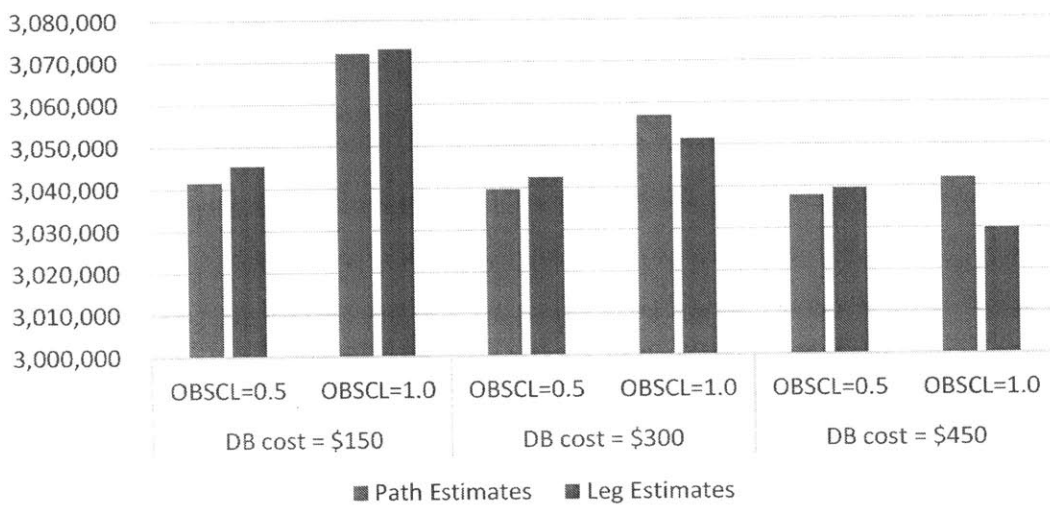

Figure 40: Test 4- Net Revenues

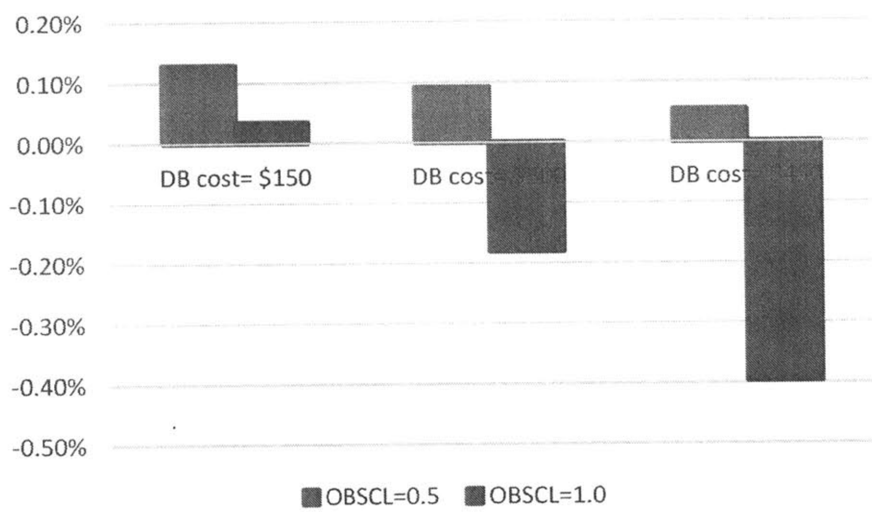

Figure 41: Test 4- Leg Estimates Revenue Gains over Path Estimates

Figure 42 and Figure 43 present AL1's load factors and yields with both approaches. Overall, there are minor differences in the load factors and yields of both approaches. AL1 has a slightly higher load factor with the leg estimates compared with the path estimates when OBSCL=0.5 and lower load factor when OBSCL=1. AL1 yields are approximately 0.15% lower with the leg estimates than with the path estimates.

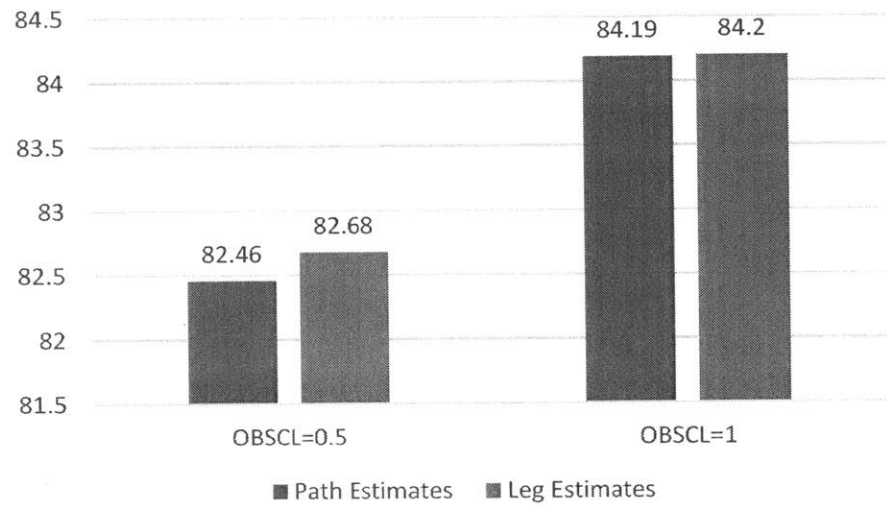

Figure 42: Test 4- Load Factors

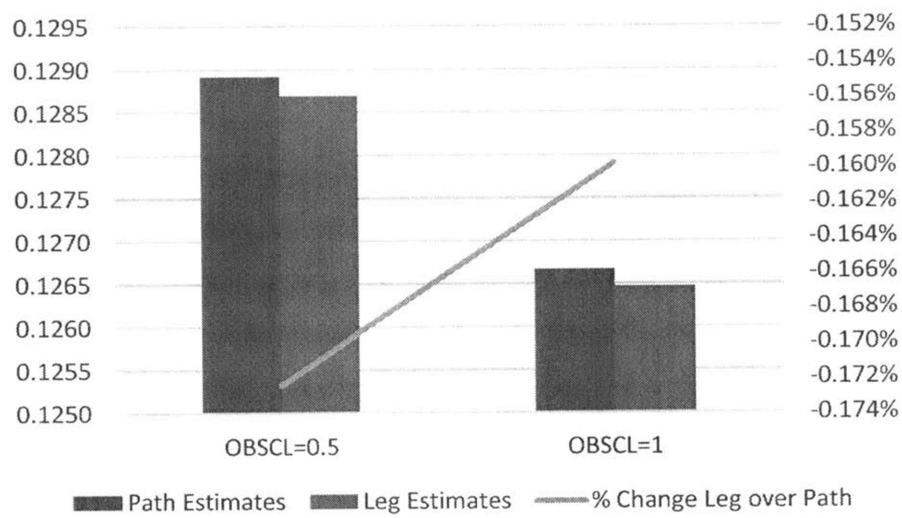

Figure 43: Test 4- Yields

The difference in total AL1 net bookings presented in Figure 44 explains the difference in load factors and yields of the two aggregation approaches. AL1 has a few more bookings in CL6, CL9 and CL10 with the leg estimates approach than with the other approach. For this reason AL1 has more DBs when using leg CXL estimates.

A separate study on the accuracy of the cancellation forecasting methods showed that, in general, the path level cancellation forecasting methods tend to slightly underestimate cancellation rates. In other words, the forecasted cancellation rates were slightly lower than the actual cancellation rates. The study also compared the accuracy of the leg level versus the path level cancellation forecasting by comparing the mean average percent errors of both levels by class and time frame. The results showed that, overall, the leg level cancellation estimates were more accurate. These findings can explain the reason for the difference in DBs between the two aggregation approaches. Since the leg level approach is more accurate and does not underestimate cancellation rates as much, it forecasts more cancellations and hence sets higher AUs than the other approach. Higher AUs translates to more bookings in the systems, which can lead to higher DBs.

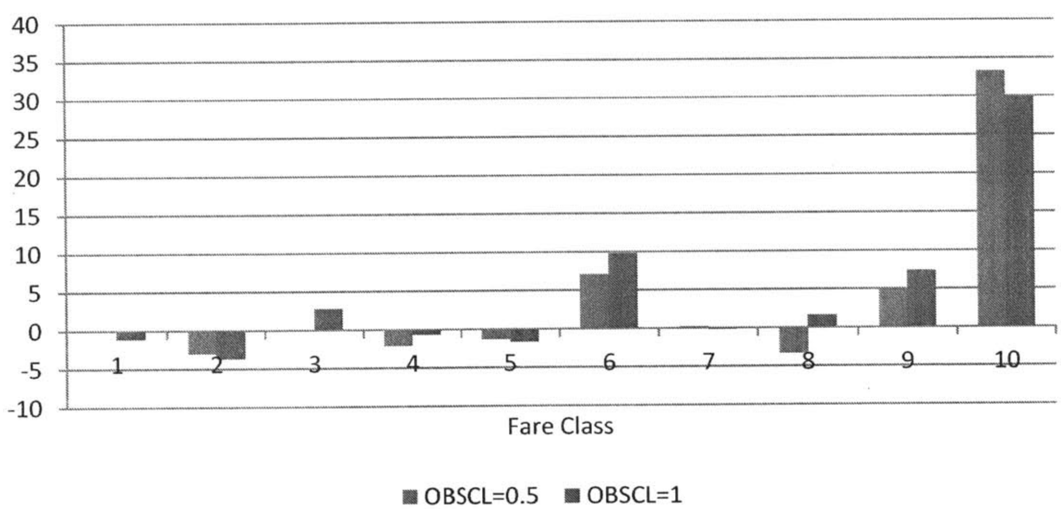

Figure 44: Test 4- Leg over Path Net Bookings by Class

Based on the results of this test, changing the aggregation level of the cancellation rates does not have a substantial effect on the results overall. Ticket revenues are indeed higher with the leg estimates, however the DBs are higher as well. The bottom line net revenues of the leg estimates approach are higher if the overbooking is moderate and lower if the overbooking is more aggressive. The fare class mix of AL1 is worse off when leg estimates approach is used as AL1 accepts slightly more low fare class bookings, which lead to marginally higher load factors and lower yields.

4.6 Test 5- Apex Overbooking (APOB)

In the fifth test, the ("new") apex overbooking methodology will be tested and compared with the ("old") OBSCL overbooking methodology which was used so far in all tests. The apex overbooking methodology was developed in order to limit the number of DBs in each scenario and with it mitigate the negative effect on airlines' revenues. The settings for this test are as follows:

- Medium demand level (such that load factor will be approximately 82%)
- Medium CXL Rate (with the same probabilities presented in Figure 6.1)
- Optimizer: AL1- DAVN, AL2- DAVN, AL3- EMSRb, AL4- DAVN
- Forecaster: All airlines use standard forecast
- AL1 uses CM3 and OBSCL=0.5, 1.0, 1.5
- Other airlines use CM2 and OBSCL=0.5

Figures 45 through 47 illustrate the main difference in the two overbooking approaches by showing the average percentage overbooking over (physical) capacity by leg for each OBSCL used. When OBSCL=0.5, it is apparent that the percentage overbooking of the two methods is very similar, if not identical. As a reminder, the AU set by the new method is equal to the minimum of the AU set by the APOB technique or the AU set by the OBSCL technique. When overbooking is moderate, the AU set by the OBSCL technique appears to be the lowest in all time frames and thus there is no difference between the two methods.

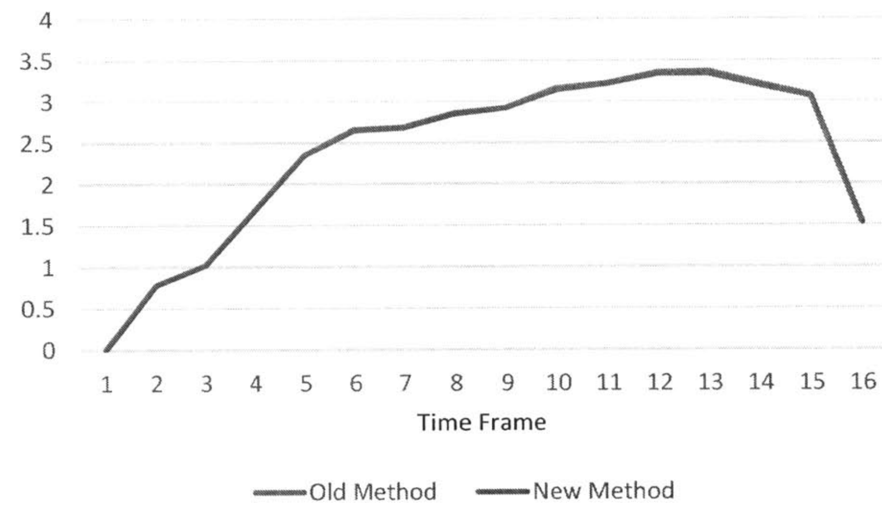

Figure 45: Test 5- % Overbooking over Capacity- OBSCL=0.5

When OBSCL=1.0, there is a clear distinction between the percentage overbooking of the two methods. The percentage overbooking is very similar with both methods up to TF4 though afterwards the old method's percentage overbooking is above that of the new method's up to departure. The gap between the two methods reaches almost two percent points gap in TF13. The curves in Figure 46 imply that up to TF4 the AU that was set by the OBSCL technique was the lower AU, while from that time frame onwards the AU set by the APOB technique was lower. It should be noted that the percentage overbooking in the OBSCL=1.0 scenario is higher than in the OBSCL=0.5 scenario in Figure 45. When OBSCL=1.5, the gap between the percentage overbooking curves is even larger. The percentage overbooking of the old method is significantly higher starting in TF3. Nonetheless, the percentage overbooking of the new method in Figure 47 is similar to that in Figure 46, as the apex of the historical BIH observations in both cases is the same. For this reason AL1's results are expected to be similar in both scenarios.

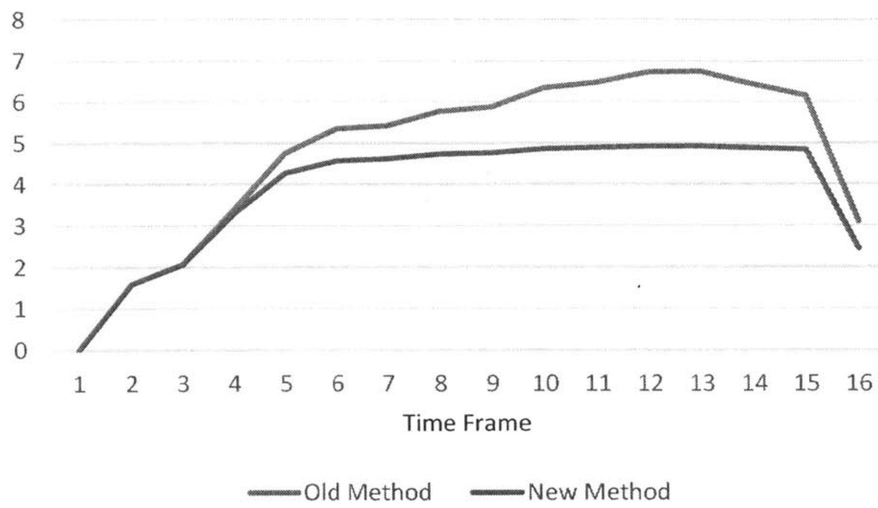

Figure 46: Test 5- % Overbooking over Capacity- OBSCL=1.0

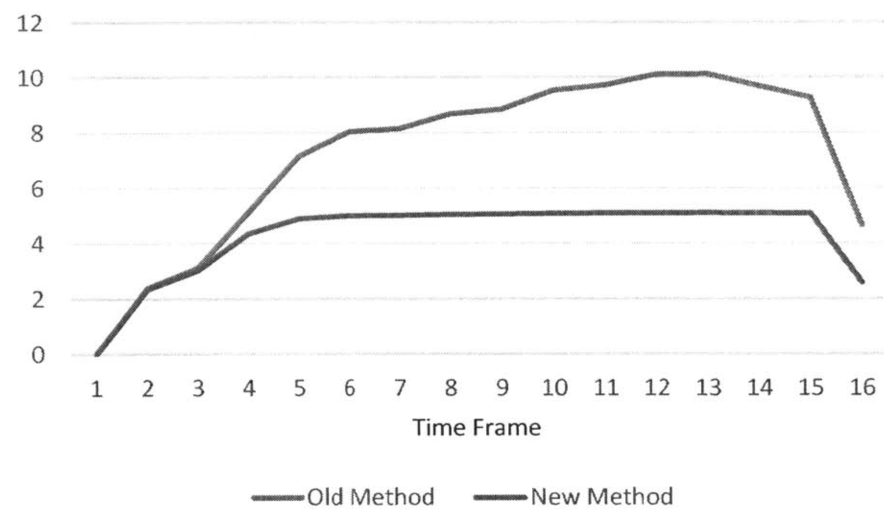

Figure 47: Test 5- % Overbooking over Capacity- OBSCL=1.5

The ticket revenues of AL1 with both overbooking methods are shown in Figure 48. Overall, the new method has a negative effect on AL1's ticket revenues. When OBSCL=0.5, there is a very marginal difference in ticket revenues. When OBSCL=1 and OBSCL=1.5 AL1's ticket revenues are 0.42% and 1.46% lower, respectively, with the new method. On the other hand, the DB numbers are significantly lower as seen in Figure 49, especially if AL1 wants to use more aggressive overbooking measures. When

OBSCL=0.5, the number of DBs is very similar with both methods, as expected. When OBSCL>0.5, there are approximately 21 DBs per 10,000 passenger booked, which is a more reasonable number, though still more than double the industry standard. Based on these numbers, it seems that the new method does fulfill its purpose of limiting the number of DBs to a level that is much closer to accepted industry levels.

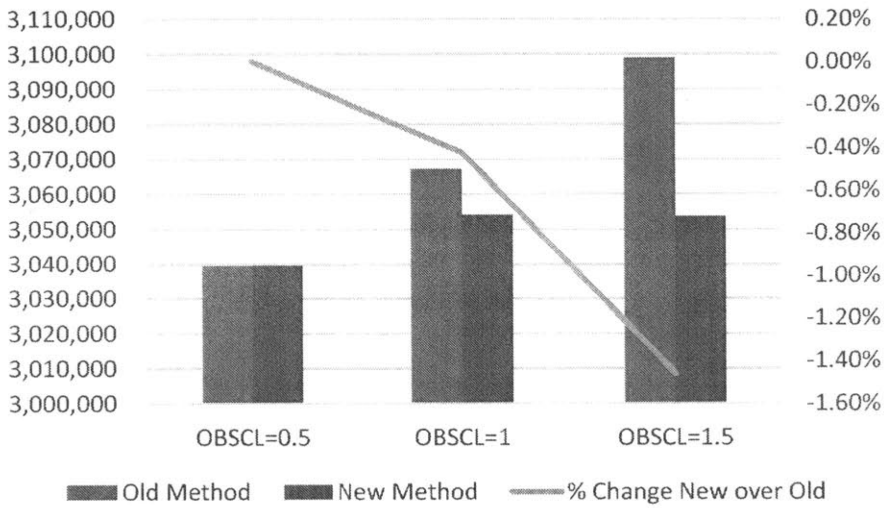

Figure 48: Test 5- Ticket Revenues

Figure 50 compares AL1's net revenues under both methods. When OBSCL=0.5, AL1 has very similar net revenues under both methods. That is due to the very similar ticket revenues and number of DBs. When OBSCL>0.5, the old method's net revenues are higher than the new method's net revenues only if the DB costs are low ($150) and lower otherwise. The net revenue gains with the new method can be up to 0.97% in the "extreme" case of OBSCL=1.5 and very high DB costs, and -0.67% when OBSCL=1.5 and DB costs are $150. Overall, the figure shows that the cap the new method sets on bookings reduces ticket revenues and also substantially mitigates the negative effect the DBs have on overall revenues.

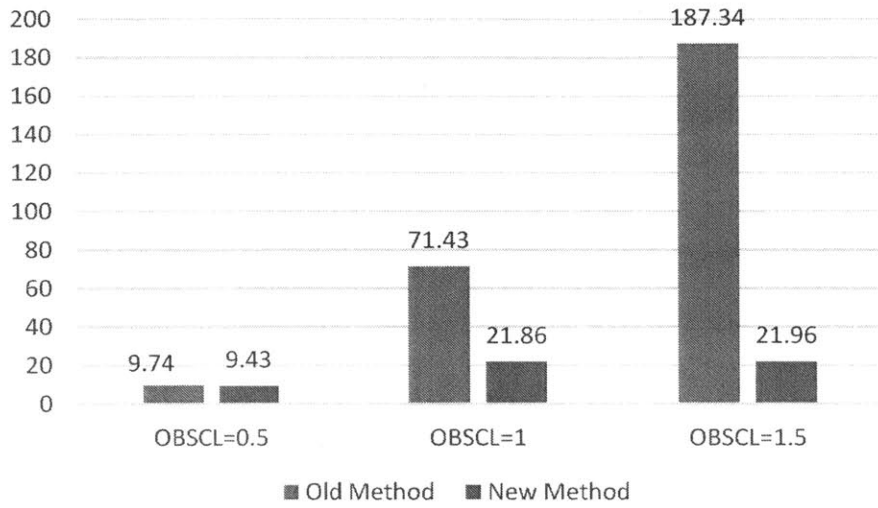

Figure 49: Test 5- Denied Boardings per 10K Passengers Booked

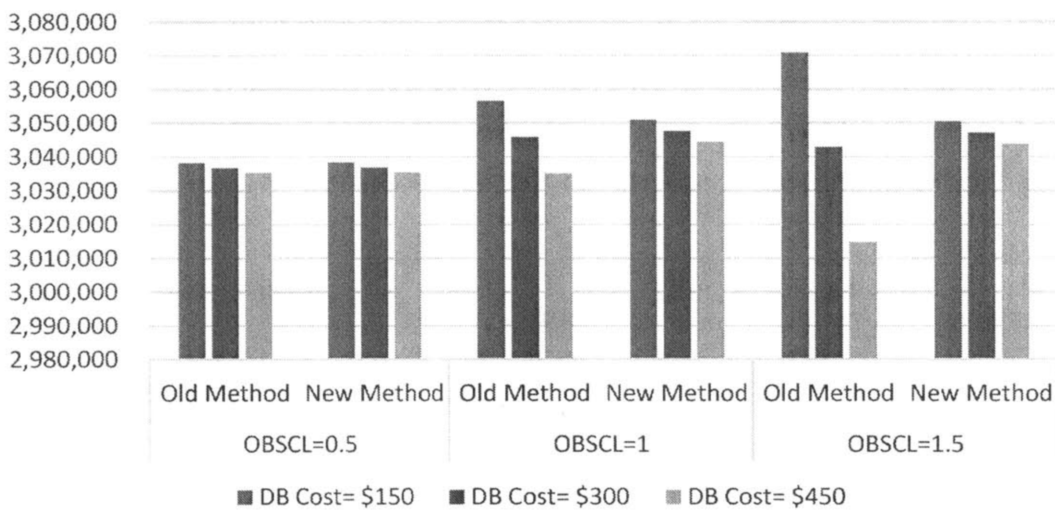

Figure 50: Test 5- Net Revenues

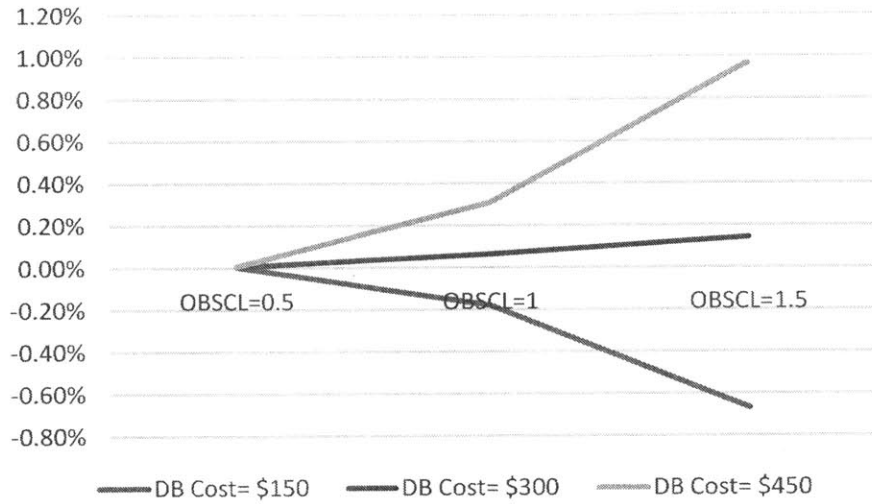

Figure 51: Test 5- % Change New Method/Old Method

Figure 52 and Figure 53 show the load factors and yields for both methods. There are very minor differences in load factors and yields when OBSCL=0.5. The load factors under the new method are about 0.3 percent and 0.4 percent points lower and, in turn, yields are 0.49% and 1.03% higher when OBSCL=1 and OBSCL=1.5, respectively.

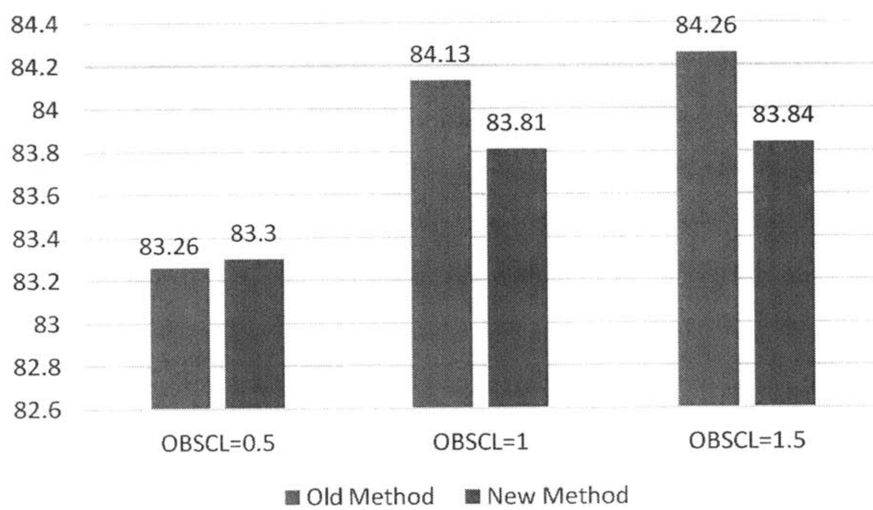

Figure 52: Test 5- Load Factors

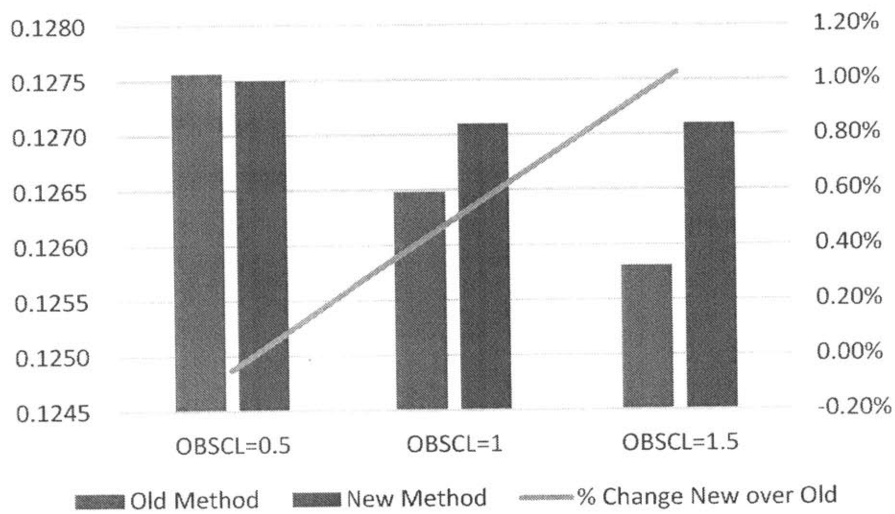

Figure 53: Test 5- Yields

The difference in load factors and yield is explained by the change in gross bookings presented in Figure 54. AL1 accepts fewer bookings under the new method, especially in CL10, due to the capping set by that method. When OBSCL=1, AL1 accepts 75 and 166 less bookings in CL10 under the new method compared with the old method. When OBSCL=0.5, the difference in bookings between the two methods is negligible. The cap the set by the new method is also clearly observed by the EMSRc presented in Figure 55. The EMSRc is not affected as much by the overbooking measures taken by AL1 as observed in Figure 21. For OBSCL=0.5, the EMSRc in both figures is similar, and for OBSCL>0.5 the EMSRc is higher with the new method. The EMSRc is similar when OBSCL=1 and OBSCL=1.5.

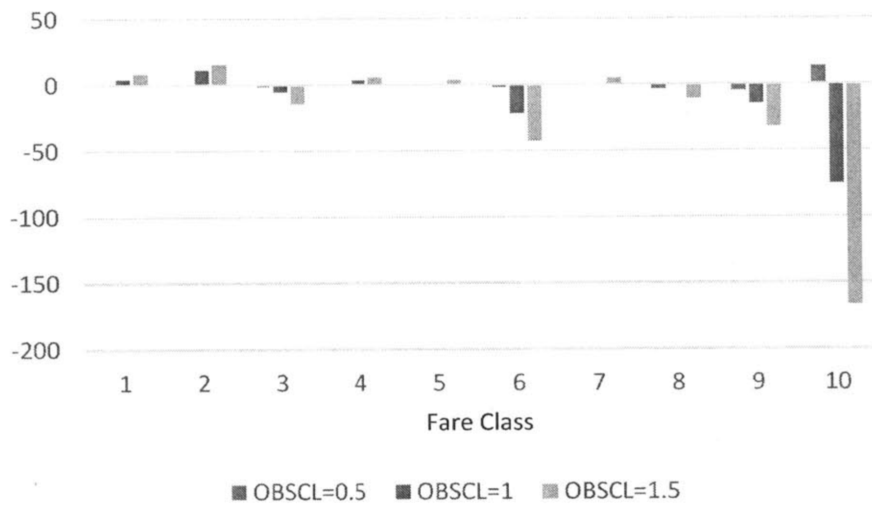

Figure 54: Test 5- New over Old Gross Bookings by Class

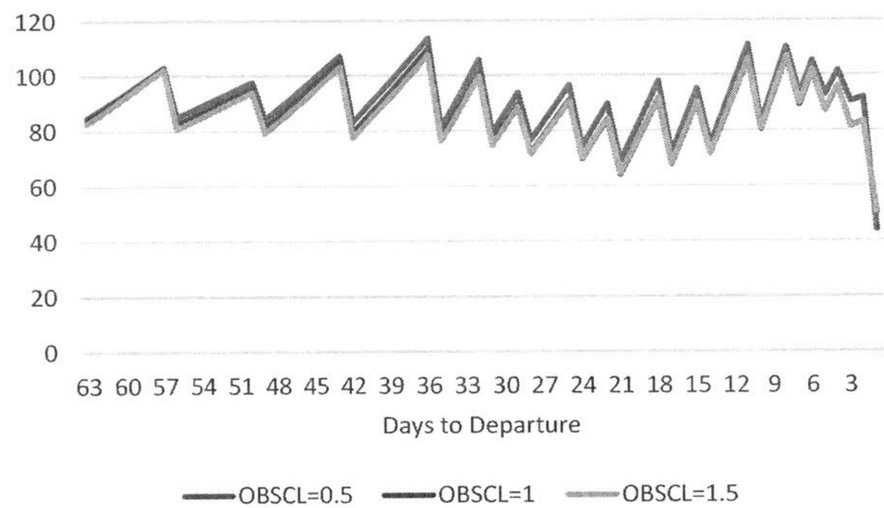

Figure 55: Test 5- EMSRc

The APOB method attempts to address the main drawback of the overbooking mechanism currently employed by the cancellation forecasting methods tested and illustrated in previous tests which is the high number of DBs when AL1 uses an OBSCL>0.5. The high number is well above the level that airlines in the real world would be willing to accept. A new overbooking mechanism was hence developed to cap the AU such that the number of DBs would be reduced. Capping the AU results in higher EMSRc, especially in early time frame which means AL1 accepts fewer lower fare class bookings compared with the old method. Consequently, ticket revenues and load factors decrease and, in turn, yields increase. Nonetheless, the goal of the new method was achieved as the number of DBs is indeed capped and closer to industry levels, though still higher. Comparing the net revenues with both methods gives a mixed result. If the DB costs are low (=$150), the net revenues are higher under the old method. The supremacy of the new method is observed only if the DB costs are higher. It is therefore implied that limiting the capacity on bookings can be beneficial for airlines only if the DB costs calculated are relatively high and might hurt bottom line revenues more than the additional gains due to overbooking. However, accurate computation of DB costs is still a challenge for airlines today.

4.7 Test 6- High Cancellation Rates

In Chapter 2 we mentioned that cancellation rates can reach 30% for some airlines due to a myriad of reasons. The sixth test in this chapter will present the performance of the three cancellation forecasting and overbooking methods in a high CXL rate environment. In this test, approximately 30% of gross bookings are cancelled before departure. For the purpose of this test, the underlying system demand had to be adjusted in order to maintain a load factor of approximately 81%. The cancellation probabilities assigned to the bookings are shown in Figure 55 (and are double the probabilities shown in Figure 6.2). The rest of the settings are as follows:

- Optimizer: AL1- DAVN, AL2- DAVN, AL3- EMSRb, AL4- DAVN
- Forecaster: All airlines use standard forecast
- AL1 uses CM1, CM3/4 and OBSCL=0, 0.5, 1.0 ("old" method)
- Other airlines use CM2 and OBSCL=0.5

Passenger Type	Penalty	TF after booking	Between	Last 2 TFs
Business	No	0.4	0.02	0.2
	Yes	0.2	0.01	0.1
Leisure	No	0.24	0.012	0.12
	Yes	0.12	0.006	0.06

Figure 56: Test 6- CXL Probabilities for High CXL Rate Scenario

Figure 56 presents AL1's ticket revenues in all scenarios. As expected AL1 gains higher ticket revenues with CM3 and all OBSCL scenarios and with CM4 when OBSCL>0 than with CM1. The higher the OBSCL is the higher the revenue gains are over CM1. CM4 performs worse than CM1 only when overbooking is not applied. Again, the RM system in the CM4 and OBSCL=0 scenario is overly restrictive and opens fewer seats in lower fare classes. Despite the similarity of these results to the results in Test 1 and Test 2, the magnitude of the revenue gains over CM1 is substantially greater in this test. However, as in previous tests, the ticket revenues do not reflect the negative effect of DBs.

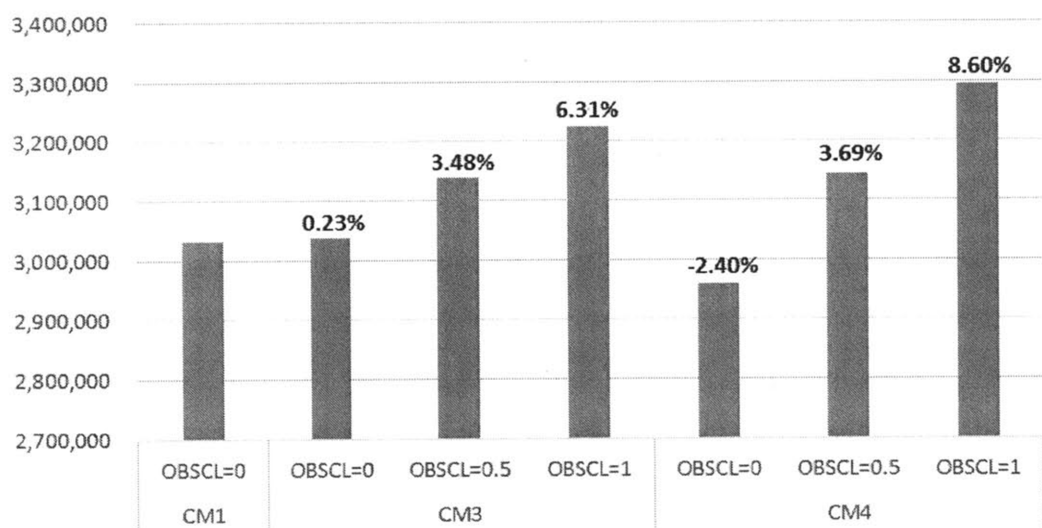

Figure 57: Test 6- Ticket Revenues

The DBs presented in Figure 58 are similar to the DBs presented in Test 2 (Figure 23) when OBSCL=0.5, however the number of DBs is noticeably higher for this test compared with Test 2 when OBSCL=1. This is a direct consequence of the overbooking methodology which assumes a linear relationship between cancellation rates and overbooking of remaining capacity. In both OBSCL scenarios, CM4 has more DBs than CM3.

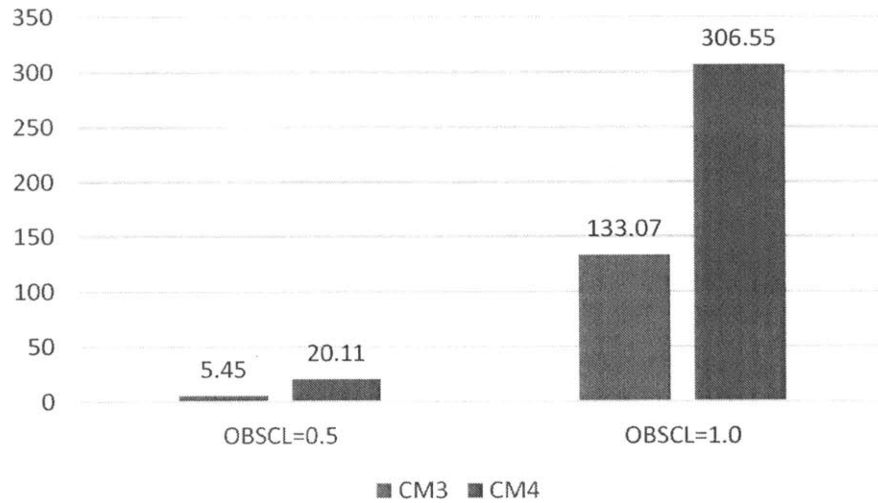

Figure 58: Test 6- Denied Boardings per 10K Passengers Booked

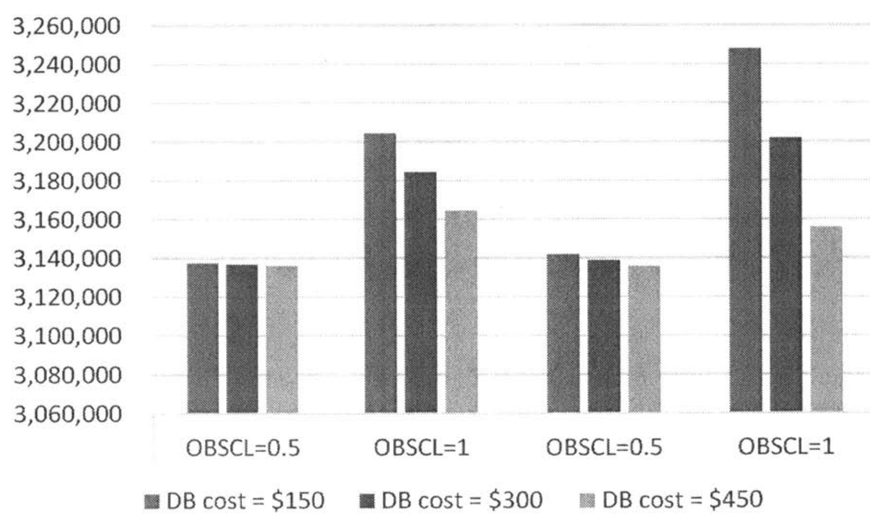

Figure 59: Test 6- Net Revenues

The net revenues of AL1 with CM3 and CM4 are presented in Figure 59, and the percent net revenue gain over CM1 is presented in Figure 60. The net revenues are higher when OBSCL=1 both with CM3 and CM4 despite the bigger number of DBs compared with previous tests. The revenue gains over CM1 are bigger as well compared with previous tests.

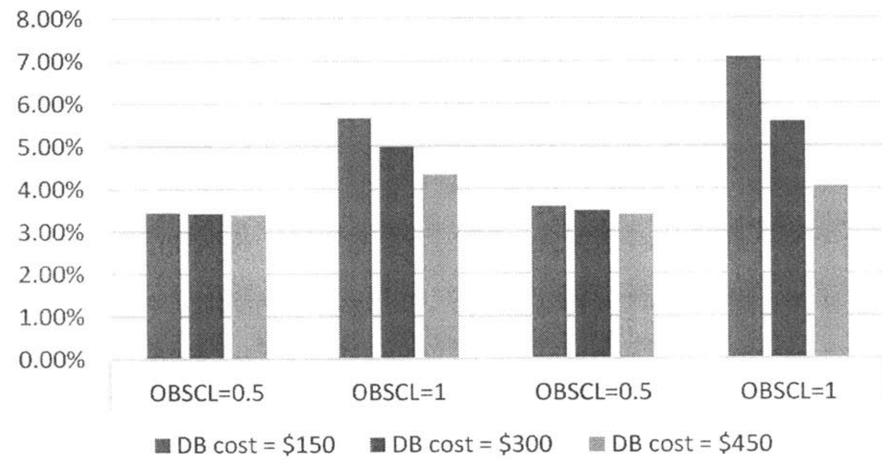

Figure 60: Test 6- % Net Revenue Gain over CM1

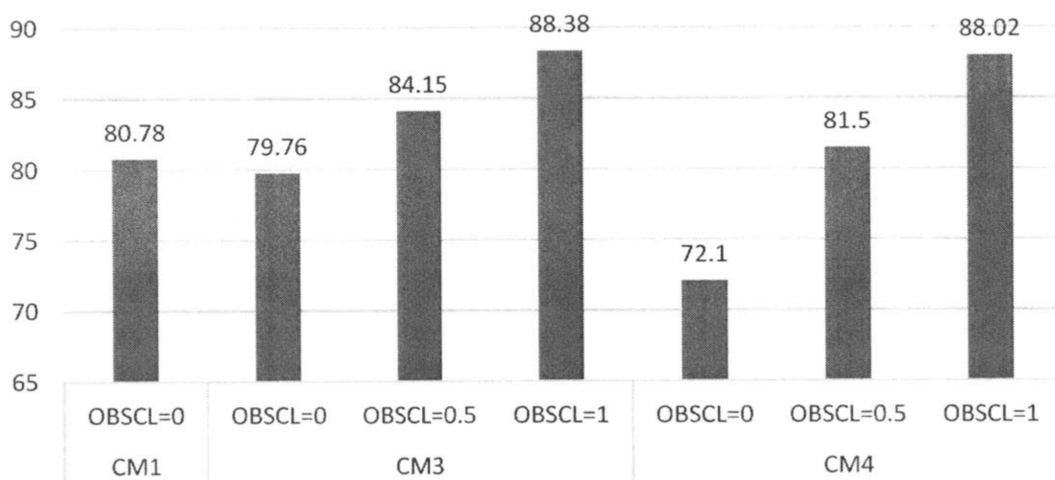

Figure 61: Test 6- Load Factors

Figures 61 and 62 present the load factors and yields for AL1. The trends in these figures are similar to the trends in previous tests: i) CM3 and CM4 have lower load factors, and in turn, higher yields than CM1 when overbooking is not applied, ii) overbooking increases load factors and decreases yields, iii) Changes in load factors and yields due to overbooking are greater with CM4. The main difference between this test and previous tests is the bigger magnitude of changes compared with CM1.

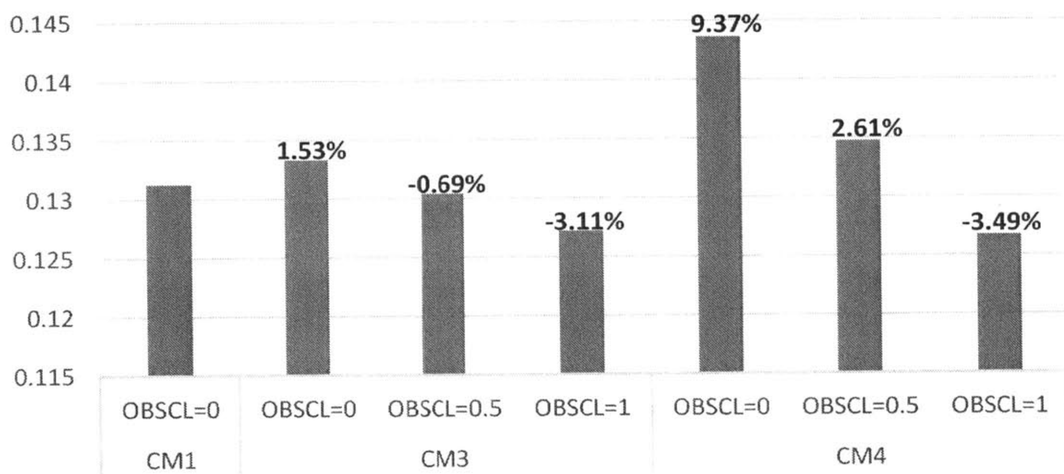

Figure 62: Test 6- Yields & Percent Gain over CM1

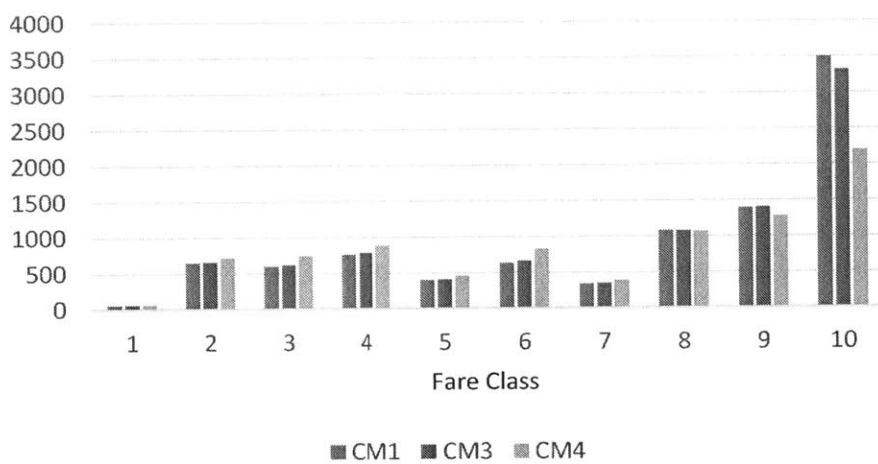

Figure 63: Test 6- Net Bookings without Overbooking

Figures 63 and 64 compare the net bookings for CM3 and CM4 without overbooking and with overbooking. The net bookings trends in both figures are similar to the trends in previous tests.

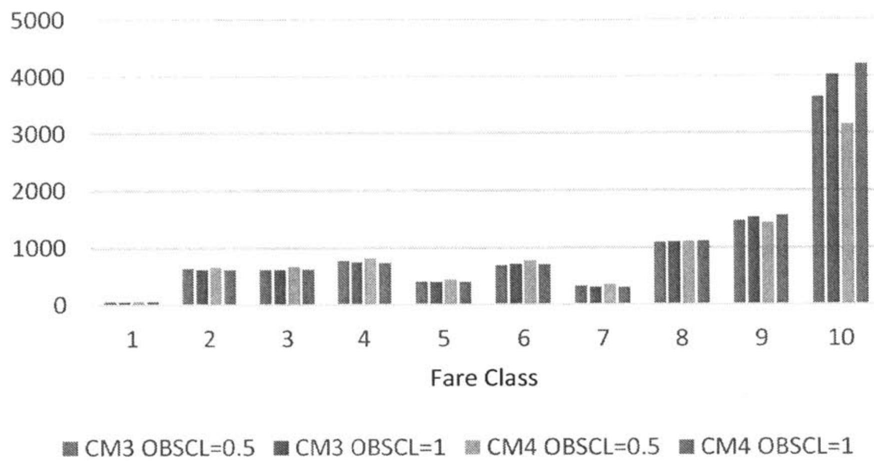

Figure 64: Test 6- Net Bookings with Overbooking

To summarize, this test analyzed the effects of cancellation forecasting and overbooking in a high cancellation rate scenario. The trends of the results in this test are similar to the trends in Tests 1 and 2. The main difference between the tests is the magnitude, both in absolute and percentage terms, of the revenue gains from cancellation forecasting and overbooking. In this test, the revenue gains over CM1 (or no cancellation forecasting) in some cases are more than double the revenue gains in previous tests, therefore making the advantages of cancellation forecasting and overbooking even more apparent.

4.8 Test 7- RM optimizers

The seventh test compares the performance of three different RM optimizers, taking into account passenger cancellations. Some of the previous runs and theses using PODS have analyzed the performance of different optimizers under different scenarios, however passenger cancellations were not considered. Thus, this test includes three scenarios and in each scenario AL1 uses one of the three optimizers discussed earlier in this chapter: DAVN, ProBP and UDP. In addition, the test will also address the difference in the RM optimizers under two cancellation forecasting methods, CM3 and CM4, and three different overbooking levels, OBSCL=0, 0.5 and 1. AL1's performance parameters will be presented and compared between the scenarios. The settings for this test are as follows:

- Medium demand level (such that load factor is approximately 82%)
- Medium CXL Rate (with the same probabilities presented in Figure 6.1)
- Optimizer: AL2- DAVN, AL3- EMSRb, AL4- DAVN
- Forecaster: All airlines use standard forecast
- AL2, AL3 and AL4 use CM2 with OBSCL=0.5 (moderate overbooking)

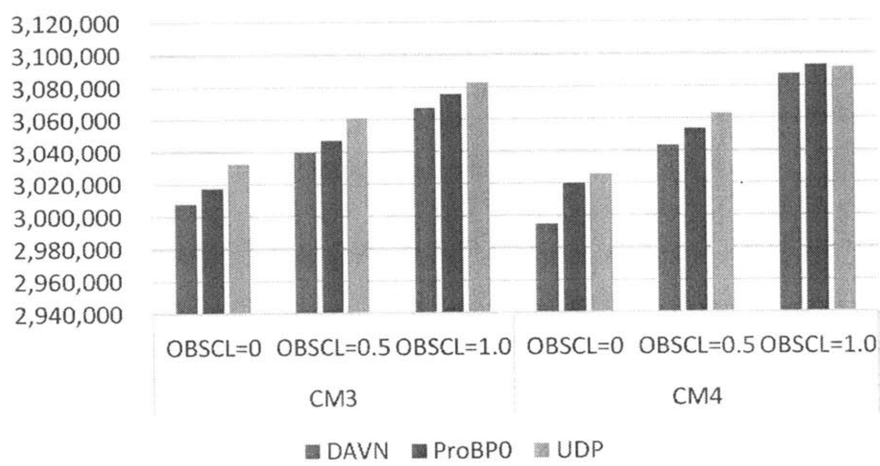

Figure 65: Test 7- Ticket Revenues

Figure 65 presents AL1 ticket revenues for all scenarios. In all scenarios AL1 gains the highest ticket revenues with UDP, followed by ProBP and DAVN gains the lowest ticket revenues. When OBSCL=0, AL1 has higher revenue gains with CM3 than with CM4 when using DAVN and UDP, which is in line with results in test 2, and lower ticket revenues when using ProBP. When OBSCL>0, AL1 has higher ticket revenues with CM4 with all optimizers. Nonetheless, as seen in test 2 and Figure 66, CM4 has more DBs than CM3 due to the compensation of both BIH and BTC cancellations. UDP has the substantially lowest number of DBs than DAVN and ProBP in both OBSCL scenarios, though still much higher than industry standard when OBSCL=1. When OBSCL=0.5, DAVN has the highest number of DBs in the CM3 scenario, while ProBP has the highest number in the CM4 scenario. When OBSCL=1, ProBP has the highest number of DBs under both methods. Taking into account the costs associated with DBs results in the net revenues presented in Figure 67 and Figure 68.

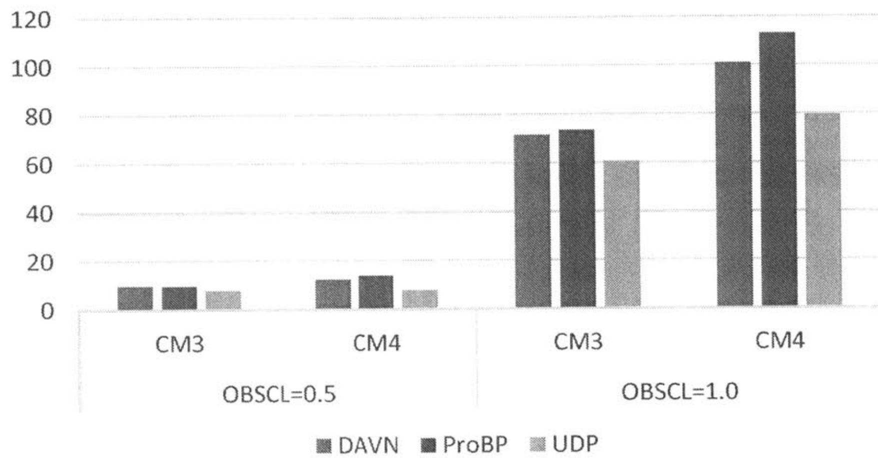

Figure 66: Test 7- Denied Boardings per 10K Passengers Booked

Similar to the ticket revenues presented earlier, AL1 still gains the highest net revenues with UDP, then with ProBP and DAVN both with CM3 and CM4, though the difference between DAVN and ProBP is marginal in the scenarios where DB costs are either $300 or $450. Overall, similar to the results in test 2, more aggressive overbooking can lead to higher net revenues if the DB costs are not too high.

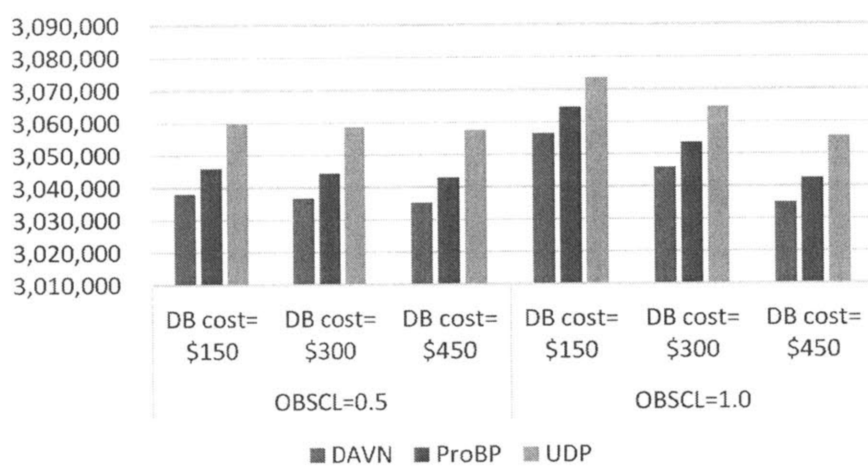

Figure 67: Test 7- Net Revenues with CM3

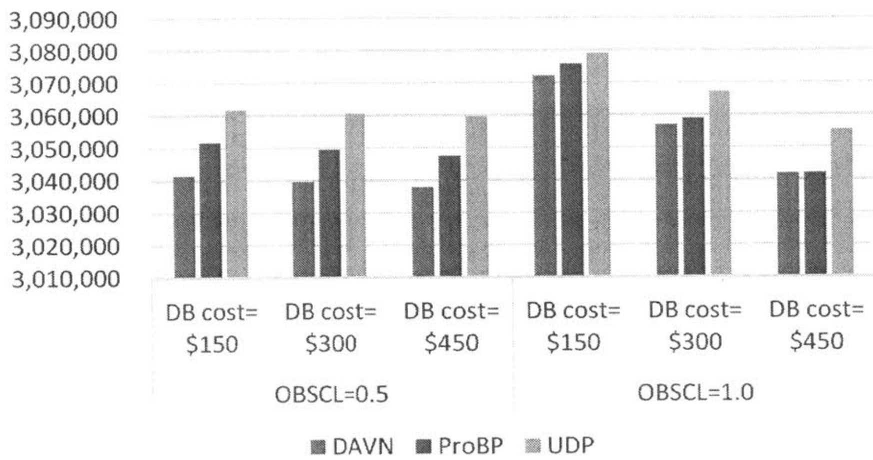

Figure 68: Test 7- Net Revenues with CM4

Figures 69 and 70 present the percent net revenue gains over CM1 when OBSCL=0.5 and 1, respectively. With CM3, the highest net revenue gains are achieved with DAVN, then with ProBP and finally with UDP. This is in contrast to the trend appearing in the figures showing the ticket and net revenues where UDP had the highest revenues, followed by ProBP and DAVN. When OBSCL=0.5, the average percent gain is approximately 1% as DB costs do not have a large effect on ticket revenues, and when OBSCL=1 the percentage gain ranges between 1.7% and 0.8%.

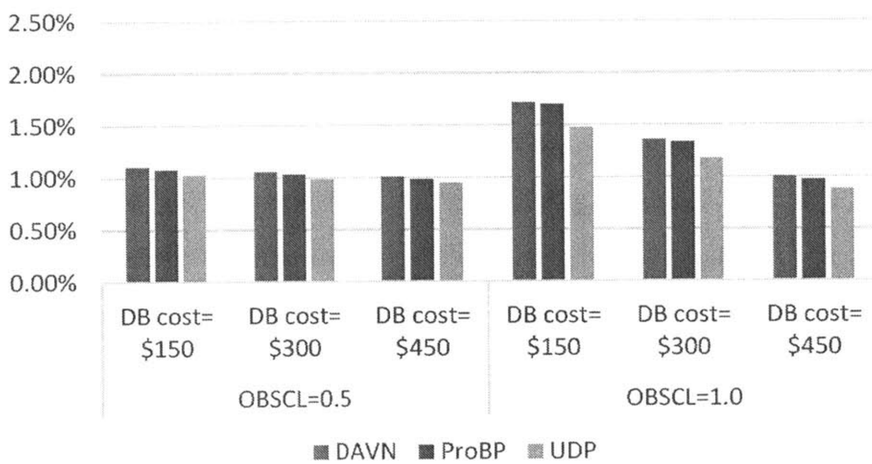

Figure 69: Test 7- CM3 % Net Revenue Gain over CM1

With CM4, the highest revenue gains are with ProBP when OBSCL=0.5 and then DAVN. The revenue gains are slightly higher than the equivalent gains under CM3. When OBSCL=1.0, the revenue gains are again higher with DAVN, then with ProBP and UDP. The DB costs have a greater impact on revenues under CM4 due to the bigger number of DBs and hence the revenue gain range between approximately 2% when DB costs are low and 0.9& when DB costs are high.

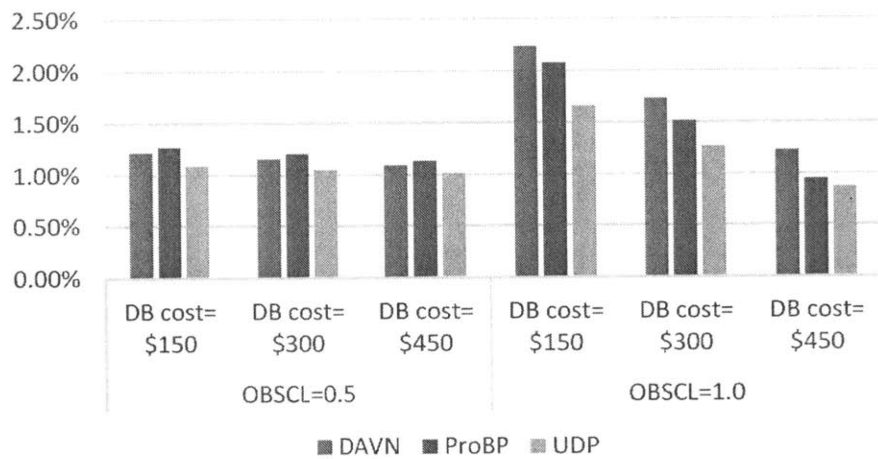

Figure 70: Test 7- CM4 % Net Revenue Gain over CM1

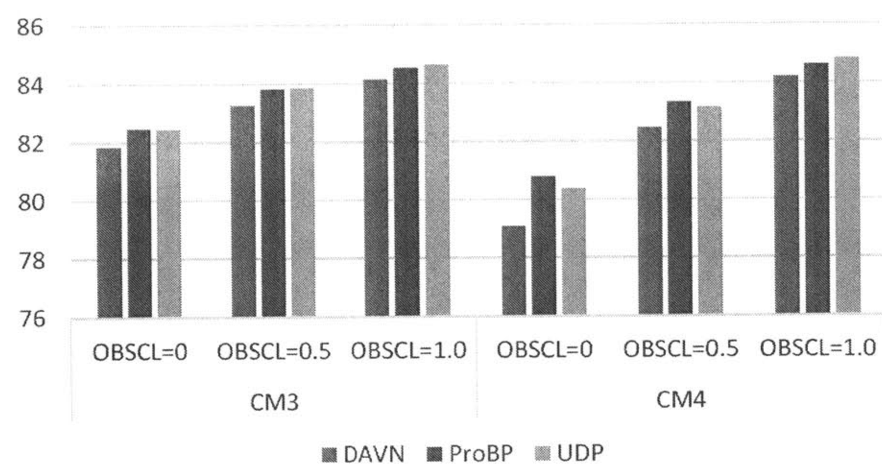

Figure 71: Test 7- Load Factors

The load factors in Figure 71 show that load factors are the lowest with DAVN in all scenarios. Under CM3, ProBP and UDP have very similar load factors while under CM4 the load factors for ProBP are higher than UDP load factors when OBSCL<1.0 and lower when OBSCL=1.0. The yields presented in Figure 72 show that in all scenarios ProBP has the lowest yields, while DAVN and UDP have higher and very similar yields under CM3. Under CM4 DAVN has higher yields. Overall the load factor and yield trends are similar to the trends shown in test 2. CM3 has higher load factors than CM4 when OBSCL<1 and similar load factors when OBSCL=1. At the same time, yields are lower for CM3 when OBSCL<1 and similar yields when OBSCL=1.

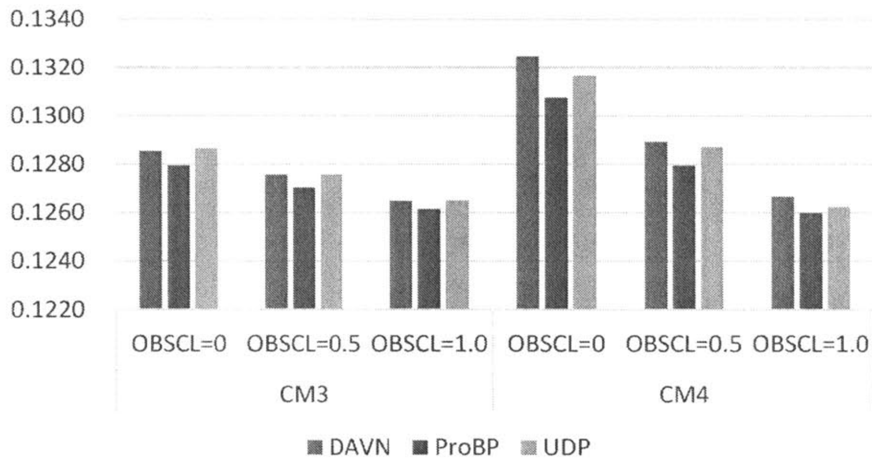

Figure 72: Test 7- Yields

Figures 73 through 75 present the difference in net bookings between ProBP and UDP over DAVN for OBSCL=0, 0.5 and 1.0, respectively. For example, Figure 73 shows that under CM3 ProBP has 35 more bookings in CL10 than DAVN while UDP has almost 10 bookings less when OBSCL=0. Under CM4, ProBP has 250 more bookings in CL10, while UDP has only 116 more bookings when OBSCL=0. For this reason there is a larger difference in the optimizers' load factors and yields under CM4 than under CM3 as seen in Figures 71 and 72. The differences in net bookings in higher fare classes are marginal.

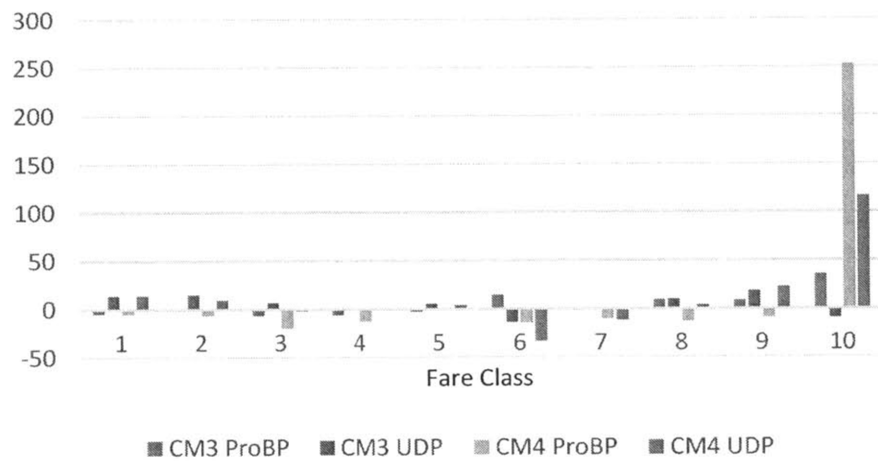

Figure 73: Test 7- Difference in Net Bookings- OBSCL=0

When OBSCL=0.5, the differences in CL10 net bookings are smaller than when OBSCL=0. Under CM3, nonetheless, ProBP has more bookings in lower fare classes and fewer bookings in higher fare classes than DAVN and hence the higher load factors and lower yields. UDP has more bookings in lower fare classes and higher fare classes and hence the higher load factor but similar yield to DAVN. Under CM4, ProBP has more bookings in lower fare classes and less bookings in high fare classes than DAVN and thus the higher load factor and lower yield than DAVN. UDP also has more bookings in CL10 than DAVN however not to the same extent as ProBP and therefore the load factor is lower and yield is higher than ProBP. The same idea repeats for OBSCL=1, however the differences in CL10 between the optimizers are even smaller than OBSCL=0 and OBSCL=0.5.

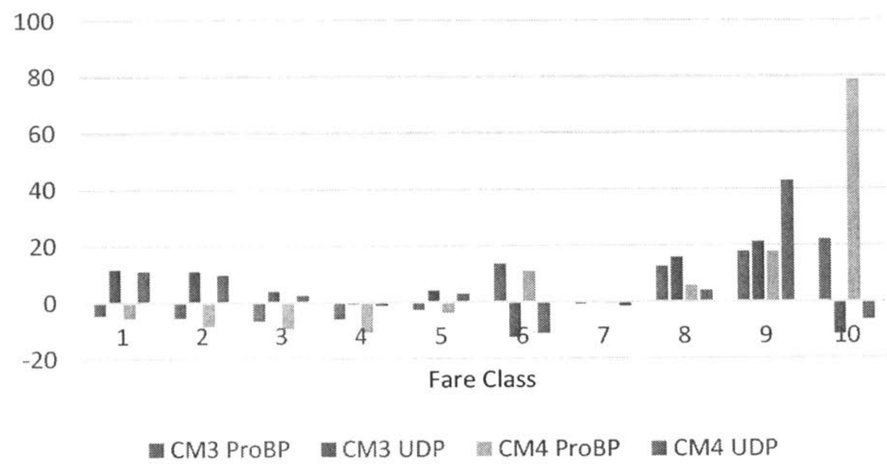

Figure 74: Test 7- Difference in Net Bookings- OBSCL=0.5

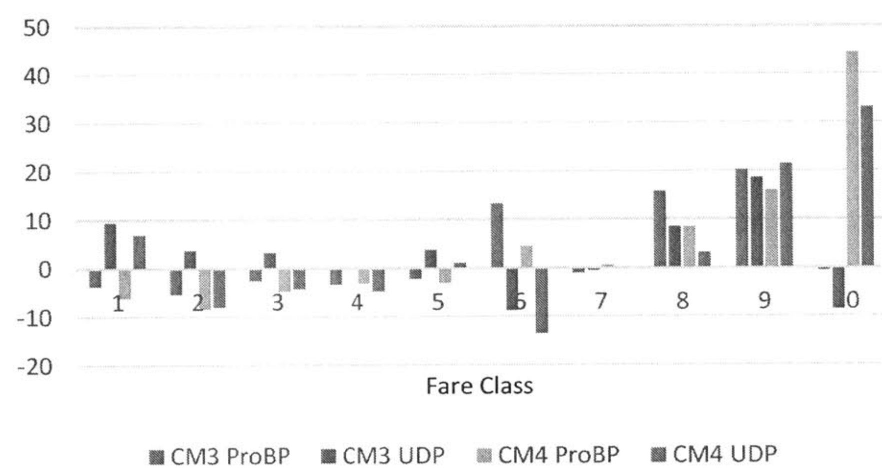

Figure 75: Test 7- Difference in Net Bookings- OBSCL=1

In conclusion, this test compared the performance of DAVN, ProBP and UDP while allowing for bookings to be cancelled during the booking process, a circumstance that makes the task of optimizing seat allocation more challenging. The results show that in all scenarios cancellation forecasting and overbooking can increase ticket and net revenues. In general, UDP is the optimizer with the highest ticket and net revenues with both CM3 and CM4, while DAVN is the optimizer with lowest revenues. The on the other hand, the revenue gains over CM1 (percentage wise) are the highest for DAVN and lowest for UDP. When OBSCL=0.5, all optimizers have a net revenue gain of approximately 1% and 1.2% over CM1 with CM3 and CM4, respectively. When OBSCL=1, the difference between the optimizers' revenue gains (over CM1) over CM1 are more apparent. The revenue gains in this case range from approximately 1% and 1.5% with CM3 and 1% and 2.2% with CM4, depending on the DB costs. UDP achieves the highest load factors, however it does not achieve the lowest yields as would be expected. The results show that UDP accepts more bookings than DAVN but mostly in higher fare classes. ProBP on the other hand accepts more bookings than DAVN but mostly in low fare classes and thus ProBP has the lowest yield between the optimizers in all scenarios. UDP also has the lowest number of DBs out of the three optimizers. This is because UDP is more restrictive during booking time frames close to departure and thus accepts fewer bookings when flight legs are booked close to or above capacity compared with the other optimizers.

4.9 Test 8- RM Forecasters

Similar to the previous test, the eighth (and last) test compares the performance of another major component in the RM system under the conditions of passenger cancellations. This thesis will be the first to present the effects of the interaction between the RM forecaster and the passenger cancellation phenomenon. Hence, this test includes three main scenarios and in each scenario AL1 uses one of the three forecaster types discussed earlier: standard forecast, hybrid forecast (HF) and hybrid forecast and fare adjustment (HF/FA). This test will also address the performance of each forecaster under CM3 and CM4 with different overbooking levels, OBSCL=0, 0.5 and 1. The rest of the settings for this test are as follows:

- Medium demand level (such that load factor is approximately 82%)
- Medium CXL Rate (with the same probabilities presented in Figure 6.1)
- Optimizer: AL1- DAVN, AL2- DAVN, AL3- EMSRb, AL4- DAVN
- Forecaster: AL2, AL3 and AL4 use standard forecast
- AL2, AL3 and AL4 use CM2 with OBSCL=0.5 (moderate overbooking)

The ticket revenues shown in Figure 76 indicate that in all scenarios SF gains the lowest ticket revenues. Under CM3, HF and HF/FA have very similar revenues when OBSCL<1, and HF/FA gains the highest revenues when OBSCL=1. Under CM4, HF gains the highest revenues when OBSCL<1, and gains similar revenues as HF/FA when OBSCL=1.

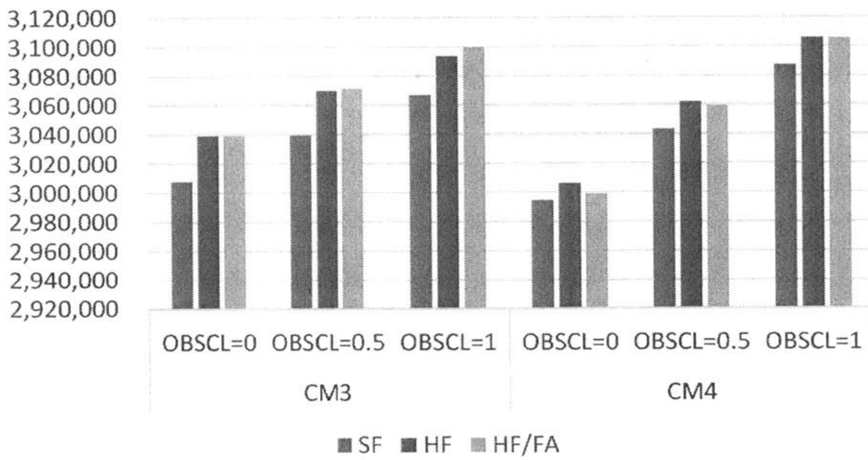

Figure 76: Test 8- Ticket Revenues

In addition the lower ticket revenues, SF has the highest number of DBs in all scenarios, followed by HF and HF/FA. The difference in DBs is high when OBSCL=1. Figure 77 shows that even with HF/FA being the most moderate in terms of DBs of 40 per 10,000 passengers it is still significantly higher than industry standard.

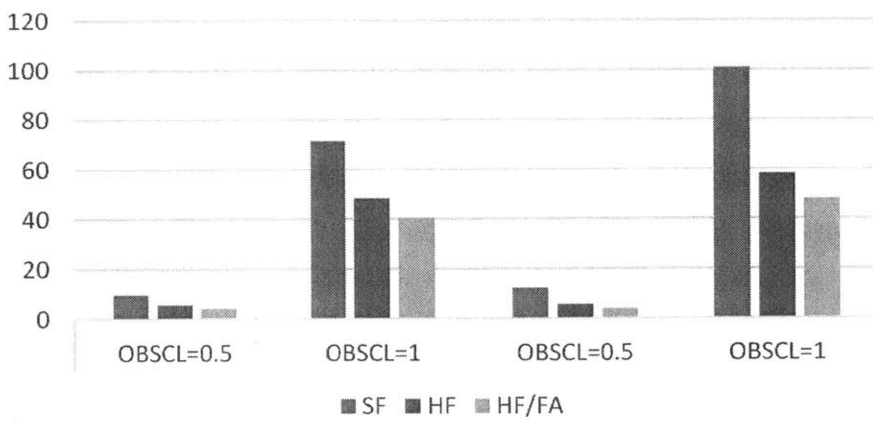

Figure 77: Test 8- Denied Boardings per 10K Passengers Booked

The net revenues under CM3 and CM4 presented in Figures 78 and 79 show that in all cases, as expected, SF has significantly lower net revenues than HF and HF/FA due to the combination of lower ticket revenues and highest DB numbers compared with the other two forecasters. The net revenues of HF and HF/FA differ according the cancellation forecasting method used. Under CM3, HF/FA has the highest net revenues in all scenarios though its advantage is more apparent when the OBSCL is higher. Under CM4, HF has higher net revenues than HF/FA when OBSCL=0.5, and lower net revenues when OBSCL=1.

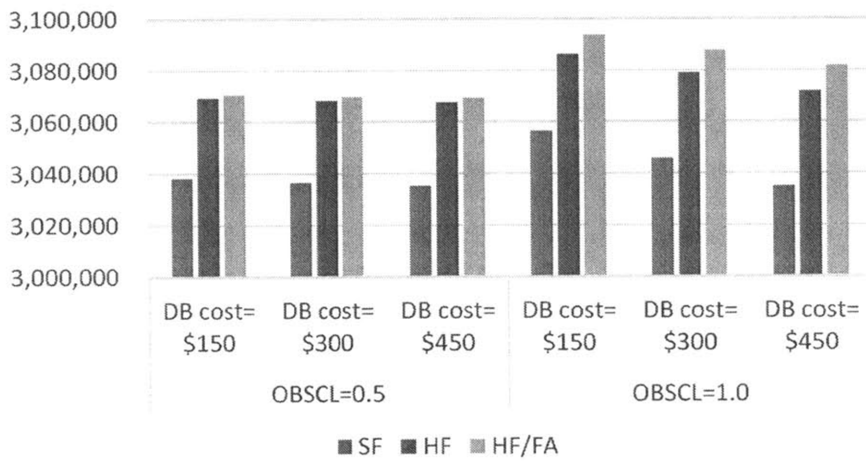

Figure 78: Test 8- CM3 Net Revenues

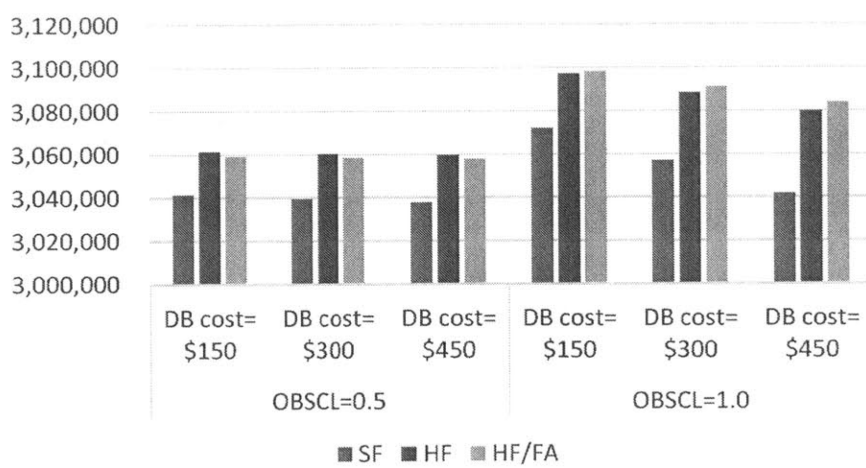

Figure 79: Test 8- CM4 Net Revenues

In order to understand the magnitude of the revenue gain with HF and HF/FA over SF, Figures 80 and 81 illustrate the revenue gain in percentages of HF and HF/FA over its equivalent SF scenario under CM3 and CM4, respectively. For example, under CM3 and OBSCL=0.5, HF gains an additional 1.03% in revenues and HF/FA gains an additional 1.07% in revenues over SF. Overall, under CM3 the revenue gains of HF and HF/FA over SF are higher and range between 1.00% and 1.5% under CM3, while under CM4 the revenue gains range between 0.6% and 1.4%, depending on the OBSCL and the DB costs. Under both CM3 and CM4, the revenue gains when OBSCL=1 are higher and difference in revenue gains between HF and HF/FA is up to 0.2 percent points.

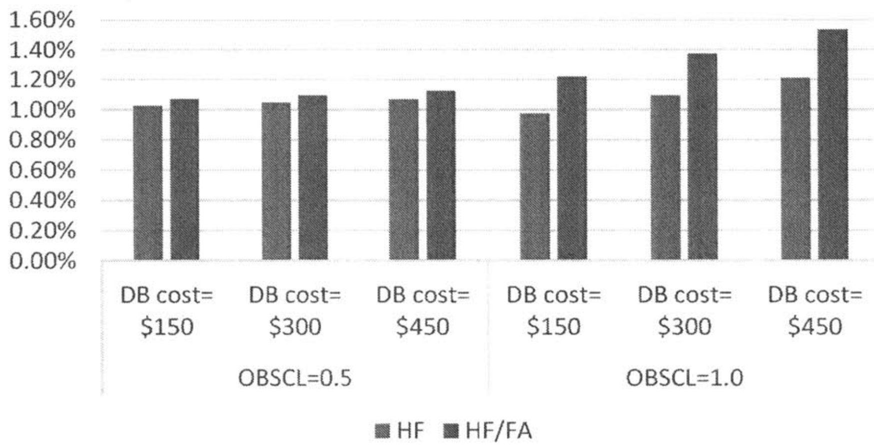

Figure 80: Test 8- CM3 Net Revenue Gain over SF

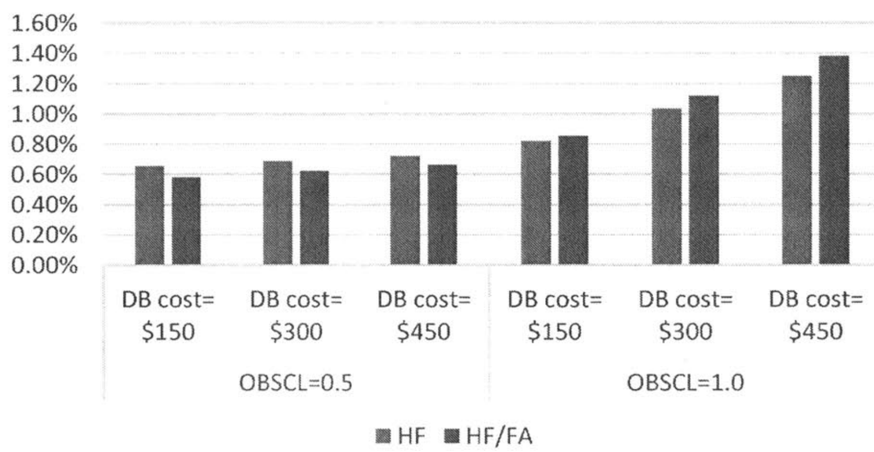

Figure 81: Test 8- CM4 Net Revenue Gain over SF

Figure 82 shows that in all scenarios the load factors are the highest with SF, followed by HF and eventually HF/FA has the lowest load factors in all scenarios. The difference in the forecasters' load factors is bigger under CM3 versus CM4. The yields in Figure 83 show that HF/FA has the highest yields, followed by HF and finally SF. The load factors and yield patterns of CM3 and CM4 are similar to the trends seen in previous tests.

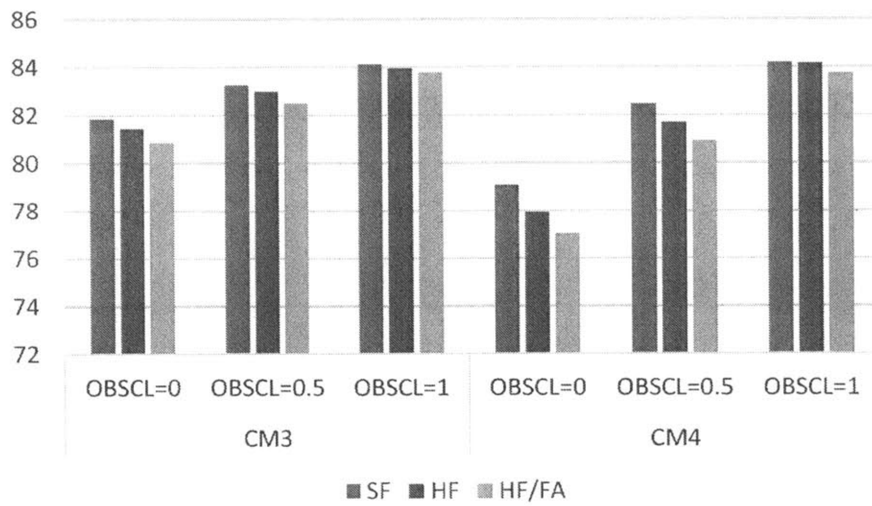

Figure 82: Test 8- Load Factors

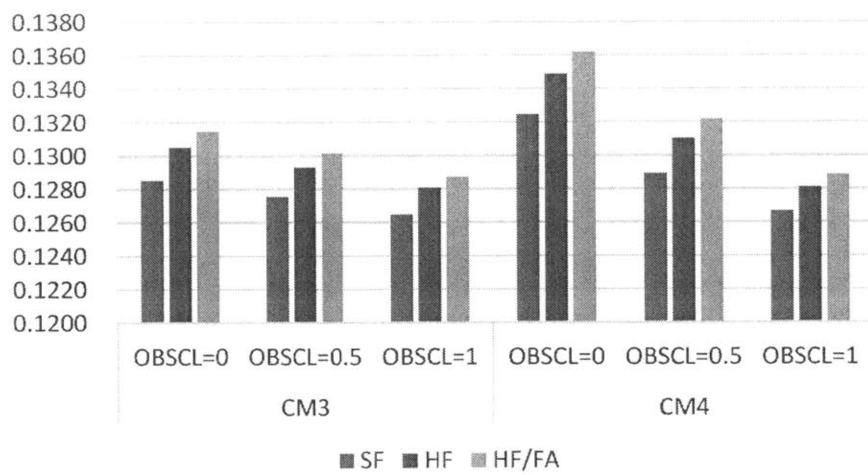

Figure 83: Test 8- Yields

The difference in the fare class mix of HF and HF/FA compared with SF when OBSCL=0.5 is presented in Figure 84. It should be noted that the results in this figure are very similar to the results when OBSCL=0 and OBSCL=1. In all scenarios, transitioning from SF to HF or HF/FA will results in less bookings in the lower fare classes, particularly CL9 and CL10 and more bookings in the higher fare classes, CL1 through CL5 and CL7. Overall, SF gains the highest number of net bookings, followed by HF and HF/FA. Both under CM3 and CM4, the difference in net bookings is bigger under HF/FA than under HF.

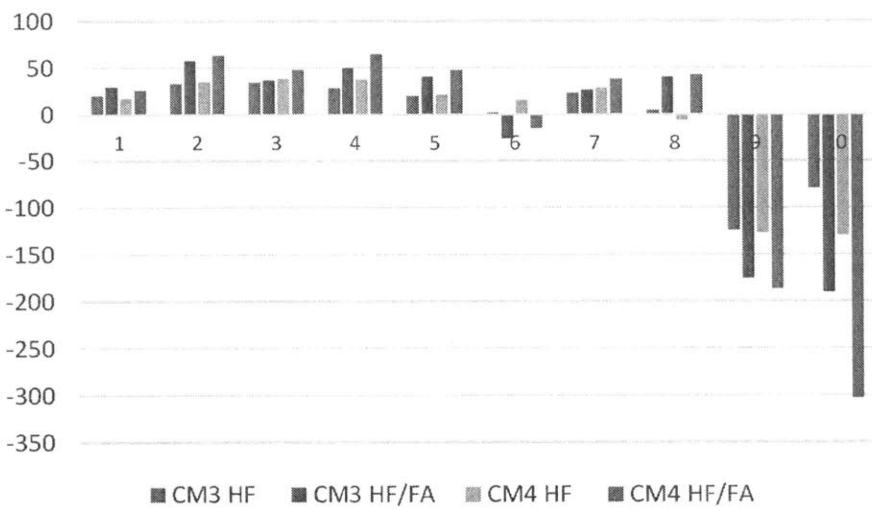

Figure 84: Test 8- Difference in Net Bookings- OBSCL=0.5

In conclusion, this test compared the performance of SF, HF and HF/FA while taking into account passenger cancellations. Previous PODS tests explored the differences in the performance of both forecasters based on solely on the methodology of each forecaster as passenger cancellations did not occur. In this section, the forecasters were compared also based on their ability to adjust for cancellations as practiced under the CM3 implementation. Under CM4 the BTC forecasts do not need to be adjusted for cancellations. The main findings in this test are:

- The ticket and net revenues are higher with HF and HF/FA than with SF in all scenarios, consistent with many previous studies.

- The net revenue gains over SF are different for CM3 and CM4. With CM3 the revenue gains range approximately 1% when OBSCL=0.5 or between 1% and 1.5% when OBSCL is higher. With CM4 the revenue gains are approximately 0.6% when OBSCL=0.5 or between 0.8% and 1.4% when OBSCL is higher
- The revenue gains over SF are higher for HF/FA with CM3 and with CM4 and OBSCL=1. The difference in revenue gains is biggest with CM3 and OBSCL=1
- HF alone performs better than HF/FA only with CM4 and OBSCL=0.5
- HF/FA performs better than HF and SF in terms of DBs
- Both HF/FA and HF have a better fare class mix as both forecasters result in less bookings in low fare classes and instead have more bookings in higher fare classes. Overall, both HF and HF/FA have slightly less bookings than SF
- Thus, HF and HF/FA have lower load factors but higher yields compared with SF

4.10 Summary

This chapter reported the findings from a series of simulation tests. The first two tests compared the performance of the four cancellation forecasting methods in PODS under different levels of overbooking, the third test compared two different competitive environments, the fourth compared two aggregation levels of cancellation estimates, the fifth compared two overbooking mechanisms, the sixth compared different CXL rate levels, and the seventh and eights tests compared the performance of different optimizers and forecasters with passengers cancellations.

The results in all tests were consistent. Cancellation forecasting was beneficial for AL1 compared to no forecasting at all (CM1). For CM2 and CM3, this was true even when no overbooking was applied (OBSCL=0) as AL1 gained an additional 0.1% and 0.12% in ticket revenues over CM1 in these scenarios. The ticket revenue gains can be increased even further with overbooking and reach 3%. The down side to overbooking is the high number of DBs, especially with the uncapped deterministic overbooking methodology which brings DB values to be well above industry standards. After taking account of different DB costs, the net revenue gains of CM2 or CM3 could range between 0.3% and 2%, depending on the overbooking aggressiveness and the DB costs. Overall, the results of CM2 and CM3 were very similar despite the different approach for the calculation of cancellation rates. The net revenue gains over CM1 could get even higher with CM4, as CM4 could gain additional 0.6% in net revenue over CM2/3 (depending on the DB costs and OBSCL) and despite the substantially higher number of DBs when OBSCL>0.5. For an airline with high cancellation rates, the importance of cancellation forecasting and overbooking is higher. In this scenario, the net revenue gains over CM1 could get up 6% with CM2/3 and 7% with CM4.

Using different optimizers and forecasters with different cancellation forecasting methods showed that some optimizers and forecasters perform better than others when cancellations are included. UDP gains the highest absolute net revenues, though the highest net revenue gains over CM1 were achieved with DAVN. The net revenue gains differ depending on the forecasting method used, the DB costs and OBSCL used. Among the forecasters HF had higher net revenues than standard forecasting. Fare adjustment increased net revenues even more in most cases. Switching to cancellation rates on a leg level has a marginal impact on ticket revenues, yield and load factors but greater impact on net revenues as leg level cancellation rates results in more DBs.

Use of more sophisticated cancellation forecasting methods and overbooking by competitors can hurt AL1 revenues due to a combination of spill (high fare class passengers book with competitors instead) and spiral down (high fare class passengers book in lower fare classes). The result for AL1 is therefore lower load factors and yields and between 0.4% and 0.8% loss of revenues. Apex overbooking is a good alternative for the OBSCL overbooking methodology as it caps the AU and thus the number of DBs to approximately 20 per 10,000, which is much closer to (though still more than double) industry acceptable levels.

Chapter 5: Use of Detailed Data in Cancellation Forecasting

As mentioned in Chapter 2, time series analysis has been the traditional approach for demand forecasting in general, and passenger cancellation forecasting in particular. The advancement in technology, combined with the fast growing data storage and processing capabilities of computers, has pushed airlines (and other service companies) to use their historical data more extensively and develop more accurate tools for forecasting purposes. Thus, in addition to using historical booking data on an aggregate level to forecast demand and cancellations for future flights, airlines have started using the more detailed information contained in their large Passenger Name Records (PNR) databases to forecast cancellation rates. The airlines use the attributes contained in the PNR data of each booking to potentially better forecast the probability of a booking in the future to cancel before the flight.

The purpose of this chapter is to explore some of the attributes included in the PNR database of an airline and gain insights whether or not the attribute in question could be significant for forecasting cancellation rates using methods such as logistic regression and others, as suggested in Romero Morales et al. (2010). Since the bookings in the current implementation of PODS do not contain PNR attributes, the analysis was based on actual PNR dataset from a North American (NA) airline. The dataset included approximately 50 million PNRs with at least one flight leg (coupon) between May 9th 2014 and May 9th 2015. The dataset had approximately 30 data elements containing information as booking time, flight departure and arrival times, Frequent Flier Program membership, passenger name, origin and destination, booking agent and point-of-sale, to name a few. The data element that contained the cancellation time (date) was only used in this analysis as an indicator whether a specific booking had been cancelled. For the sake of simplicity we assumed that all cancellations in our data were made by the passenger himself and not due to operational circumstances that forced the passenger to deviate from his original planes. Before doing the analysis, the data was filtered such that all the bookings included in the dataset were for flights operated by the airline only (not by code-share partners), in coach cabin (not premium/business/first), and ticketed.

The data mining and analysis was done using SQL programming language and SAS (Statistical Analysis System).

The literature provides several statistical tools to assess the robustness of independent variables in predicting the behavior of dependent variables. In this chapter the relationship between the attributes and passenger cancellation behavior will be examined by dividing each attribute to categories and calculating cancellation rates for the coupons that are included in the category. For the sake of simplicity, each attribute was analyzed independently, while ignoring the possible correlation between different attributes. The cancellation rates were calculated for each category on a specific reading day (RD) as follows:

$$\frac{Number\ of\ ACTIVE\ coupons\ that\ are\ cancelled\ AFTER\ reading\ day}{Number\ of\ ACTIVE\ coupons\ on\ reading\ day}$$

"Reading day" is equivalent to the data collection point (DCP) concept common among airline RM systems. The reading days chosen for this analysis are 270, 120, 60, 30, 14, 7, and 2 (days before departure). The analysis does not include no-show forecasting (i.e. cancellation rates of active PNR on day 1 and day 0 before departure) and hence cancellations on day 1 and day 0 before departure were ignored.

Figure 85 presents the cancellation rates for the attribute "number of passengers in itinerary" at each of the reading days mentioned above. The attribute was divided into two categories: "one passenger" and "more than one passenger". The idea here is to examine whether or not passengers who travel by themselves have different cancellation behavior and hence different cancellation rates than passengers who travel in groups of two or more. It should be noted that PNRs with more than ten passengers were excluded from the dataset. The figure shows that 16.55% of coupons that are active on day 270 before their respective departure date and that belong to passengers travelling alone cancel at some point between day 270 and departure date. The cancellation rates gradually decrease up to RD 30, slightly increase on RD 14 and then gradually decrease again to 7.43% on day 2 before departure. The cancellation rates of passengers travelling in groups of two or more is 10.46% on RD 270 and they gradually decrease to 4.17% on RD 2. The general trend of decreasing cancellation rates over time can be explained as follows:

- The number of bookings decreases the further away the RD is from the departure date, hence cancelled bookings sum up to a larger portion of total bookings (even if the absolute number of cancelled bookings stays constant)
- The further away the booking date is from the departure date, the higher the probability that it would need be cancelled (or changed) due to unforeseen circumstances (as the bath-tub shape cancellation pattern theory implies)

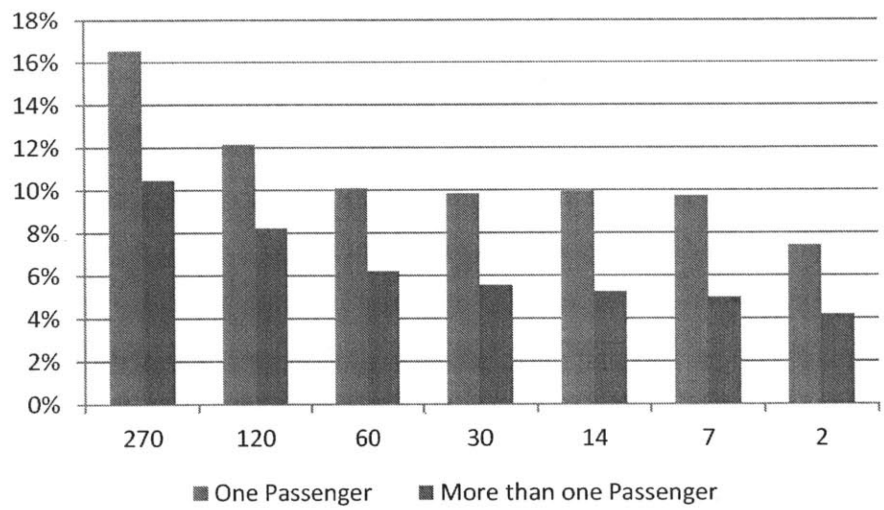

Figure 85: No. of Passengers in Itinerary CXL Rates

The second trend that stands out from Figure 85 is that, on average, bookings of passengers travelling alone have a 4.5 percent point higher cancellation rate than bookings of passengers travelling in groups of two or more. The results imply that traveling in groups of two or more reduces the motivation to cancel compared with passengers who travel alone. On the other hand, these results could also be the consequence of the fact that the two categories include different passenger types. Passengers who travel in groups are mostly leisure passengers who have firm plans and hence tend less to cancel, while passengers who travel alone could either be leisure or business passengers who tend to cancel more often.

The second attribute of interest is the "journey type" attribute. The idea is to explore whether passengers have different cancellation behavior when travelling on one way itineraries or round trip itineraries. The analysis further distinguishes between the outbound and the return portions of a round trip itinerary. The attribute hence includes three categories and their respective cancellation rates by RD are presented in Figure 86.

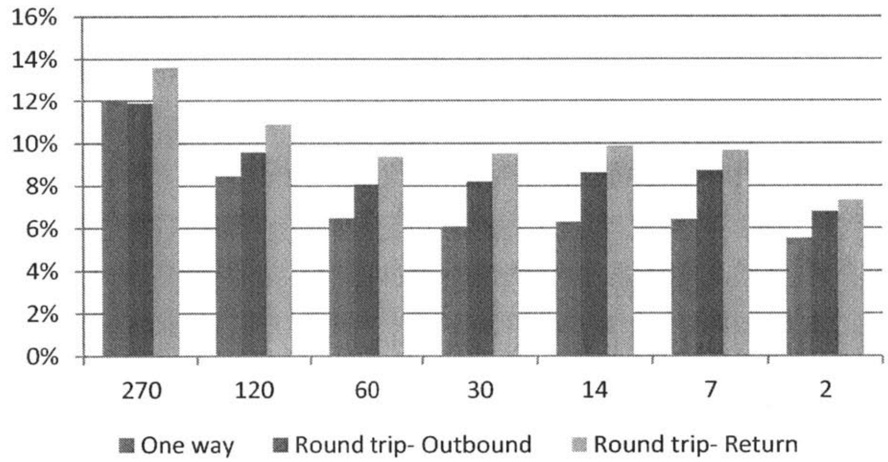

Figure 86: Journey Type CXL Rates

The relationship between cancellation rates and RDs in the figure above is similar to the relationship observed in the previous figure. That is, at first cancellation rates gradually decrease up to RD 60, then slightly increase again up to RD 14 and then gradually decrease again. This applies to all categories of this attribute. The figure also shows that bookings that consist of one way itineraries have the lowest cancellation rates (except for RD 270) while the return itinerary has the highest cancellation rates in all RDs. The difference in cancellation rates between one way itineraries and round trip itineraries could potentially be explained again by the type of passengers that purchase one way tickets versus round trip tickets. It is possible that one way itineraries have a larger percentage of leisure passengers due to their additional effort in finding one way combinations with (sometimes) different airlines. The different cancellation rates of outbound and return trips could be explained by the "bath tub" shape cancellation rates theory discussed in Chapter 3, which claims that the longer the time is between booking and departure, the higher the probability is that the passenger will cancel. Based on this theory, since the return trip is always later than the outbound trip the return trip is more subject to change. Even though it is not covered in this analysis, it would be interesting see how cancellation rates change as a function of the time between the outbound and return trips.

The third attribute that was considered in this analysis is ticket refundability. The idea was that passengers who purchase a ticket with an option to get a refund for cancelling their bookings would be more likely to cancel than passengers who do not have that option. Thus, the attribute was divided to two categories: "non-refundable" and "refundable". It should be noted that according the dataset used, only approximately 15% of tickets sold to passengers are refundable, while all other tickets are strictly non-refundable. The cancellation rates for the ticket refundability are presented in Figure 87.

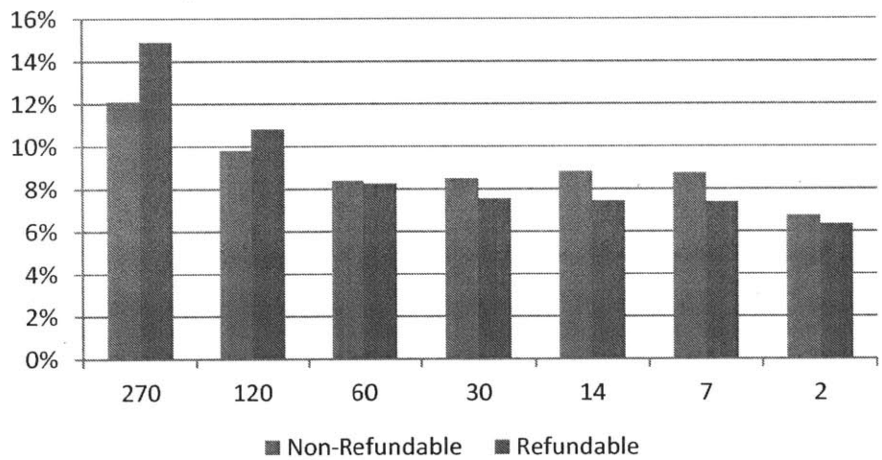

Figure 87: Ticket Refundability CXL Rates

According to the results, on RDs 270 and 120 the refundable tickets (coupons) do have a higher cancellation rate than the non-refundable tickets. On RD 270 there is almost a three percent point difference in cancellation rates and on RD 120 the difference decreases to only one percent point. On RDs 60 onwards, the refundable tickets seem to have lower cancellation rates of up to 1.5 percent points than the non-refundable tickets, a trend that is somewhat counter intuitive. Based on the results, it seems that this relatively simplistic analysis is not enough to accept or reject the hypothesis presented earlier, and thus a deeper analysis is required. It should be noted that while the total number of bookings early in booking horizon is low, a passenger who books early has more booking options (tickets with different fares and restrictions) and hence if the passenger decides to purchase a refundable ticket it could be a good indicator that he is more likely to cancel it than a purchase that purchased a non-refundable tickets. At the same time, passengers who book later in the booking horizon might purchase a refundable ticket only due to lack of other ticket options and not due to uncertainty of plans. In order to get more robust results, future analysis should keep track of the ticket types available for passengers at time of booking.

The fourth attribute is the Time of Day attribute. The idea here was to check whether or not passengers are more likely or less likely to cancel their tickets based on the time of day their flights depart. The results are presented in Figure 88.

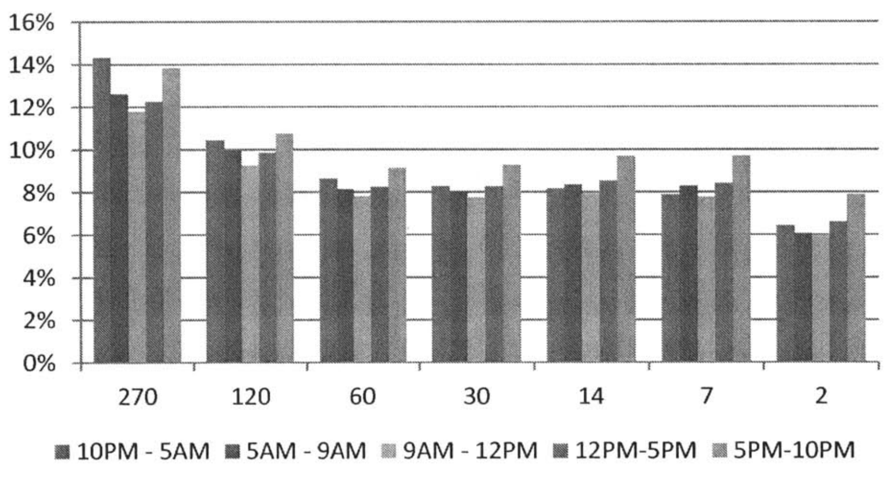

Figure 88: Time of Day CXL Rates

The results show that bookings for flight leaving late morning (9AM-12AM) have the lowest cancellation rates on all RDs. On the other hand, except for RD 270, bookings for flights leaving in the evening hours (5PM-10PM) have the highest cancellation rates. The findings could potentially be a consequence of a policy that is common among NA airlines which allows passengers to change their flights on the day of departure for a certain change fee. It is possible that passengers, especially business, who were supposed to fly later in the day decided to fly earlier due to last minute change of plans.

In addition to the day of week attribute we decided to use the day of week attribute in this analysis as well. The results are presented in Figure 89. The most noticeable trend is that weekend days, and in particular Saturdays and Sunday, have the lowest cancellation rates, while week days, in particular Tuesdays through Thursdays, have the highest cancellation rates. The trend is similar in all readings days. Just as in some of the previous attributes, the difference could be explained by the mix of passengers travelling on week days versus weekends. Business passengers travel mostly on week days and since they tend to cancel or change their plans more frequently than leisure, the cancellation rates are higher.

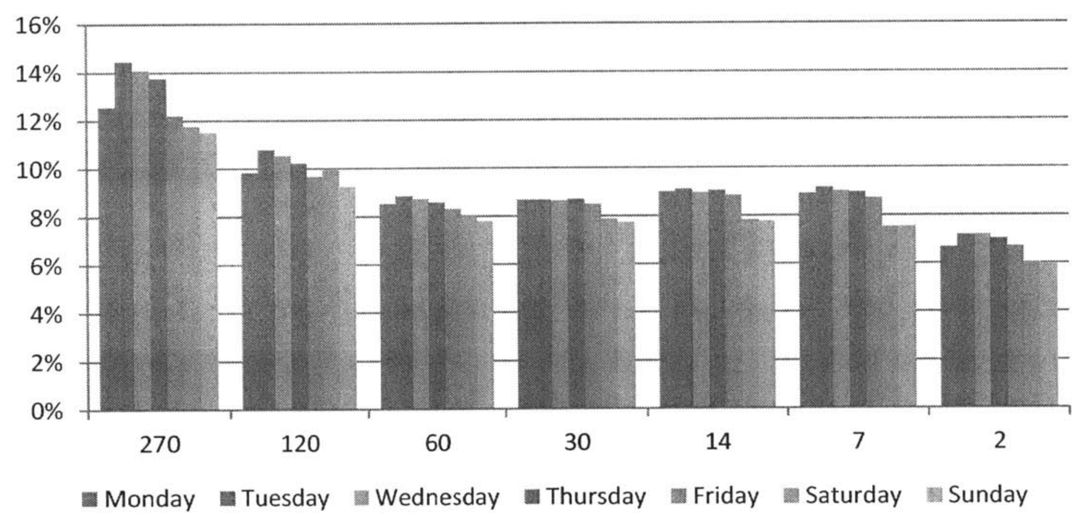

Figure 89: Day of Week CXL Rates

Another attribute that was found to be interesting for this analysis is the point of sale attribute. The idea was to examine whether there are differences in cancellation rates of booking that were made inside the US versus bookings that were made outside the US. The hypothesis is that passengers who book outside the US are mostly non-US residents with different cancellation behavior than US residents. Therefore the point of sale attribute was divided into two categories: "outside US" and "inside US". The results in Figure 90 show that at almost all RDs, the cancellation rates of the tickets booked inside the US are, on average, 1.5 percent points higher than the that tickets booked inside the US. On RD 270, the difference in cancellation rates is marginal. The results imply that non-US residents and US residents have different cancellation behavior and hence different cancellation probabilities. The reason could be either due to cultural differences or due to the additional bureaucracy non-US residents need to deal with when flying on US carriers which forces them to book only when plans are firm.

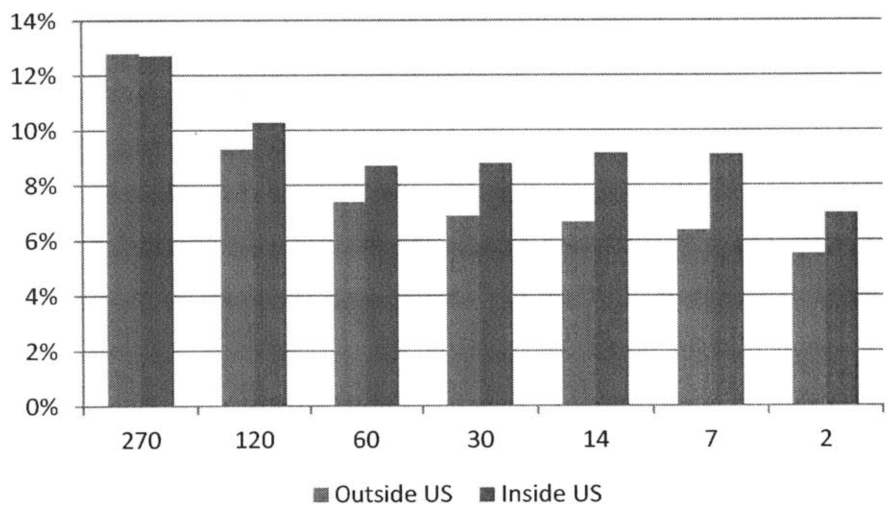

Figure 90: Point of Sale CXL Rates

In conclusion, this chapter attempted to present the importance of using disaggregate PNR data to gain better insights on passenger cancellation behavior in general and calculating cancellation rates specifically. The attributes covered in this chapter are only some of the attributes that exist in a PNR database that could be used to forecast cancellation rates at a flight leg level or OD level when used in a logistic regression. For example, the results presented in this chapter suggest that if a flight leg or OD has mostly non-US passengers, the cancellation rate could be smaller than if most passengers were US nationals, and thus the airline should be less aggressive when overbooking. Using the same logic, if a flight leg mostly consists of passengers travelling alone, the cancellation rates might be higher than if the same flight leg was mostly consisted of passengers travelling in groups, hence the airline should be more aggressive when overbooking.

The results in this chapter also show that the differences in cancellation rates could be explained by the behavior of different passenger types. Therefore, it is implied that airlines could forecast cancellation rates of flight legs or ODs based on passenger type mix. First, the attributes in the PNR can be used to classify passengers into different groups such as leisure and business, or any other type. Then, the cancellation rates of each group by RD could be calculated and used for calculating the cancellation rates of future flights. Since business passengers tend to cancel more than leisure passengers, flight legs or ODs with a large proportion of business passengers should have higher cancellation rates and hence overbooked more aggressively, and vice versa.

Chapter 6: Thesis Summary & Conclusions

Passenger cancellation (CXL) forecasting is an important element in an airline's Revenue Management (RM) system. Accurate forecasting can minimize the number of unoccupied seats at time of departure and the costs for compensating passengers who did not board due to overbooking. This thesis presented approaches for passenger cancellation forecasting and overbooking and tested them under different settings using the Passenger Origin Destination Simulator (PODS) booking simulation tool. All cancellation forecasting methods examined in this thesis are based on time series modelling of historical bookings, which is the most widespread approach by airlines with certain variations.

Four cancellation forecasting methods were presented in this thesis: CM1 through CM4. These methods differ by i) the data that is used for cancellation rate calculation, ii) the overbooking of remaining capacity, and iii) the forecast input into optimizer. CM1 is the most basic method for cancellation forecasting and thus used as a baseline for comparisons with other methods. This method does not forecast future bookings (bookings to come or 'BTC') at the start of each time frame by separating the historical data to gross bookings (the overall number of bookings in a time frame) and cancellations. Instead, the method uses historical data containing the "bottom line" net bookings (i.e. the bookings left in the system after cancellations) as a forecast. In comparison, all other methods do separate gross bookings from cancellations in each time frame. CM2 utilizes the number of cancellations in each time frame to calculate cancellation rates. CM3 utilizes the number of bookings in each time frame that would eventually be cancelled before departure. Both methods use a scaled down forecast of BTC (a forecast adjusted to cancellations) as input to the RM seat allocation optimizer.

In order to compensate for the loss of revenues due to cancellation, the remaining capacity is adjusted at the beginning of each time frame in the booking horizon to allow for more bookings than actual seats on a leg. This practice commonly used by airlines (and other industries providing perishable services) is called overbooking. The overbooking method used in CM2 and CM3 determines the overbooking level on a flight leg by calculating the product of the expected cancellations of the bookings that are currently in the system (booking in hand or 'BIH') times an overbooking scaler (OBSCL) input, by leg. The higher the expected number of cancellations or the OBSCL are, the more aggressive the overbooking will be. Since CM1 cannot forecast future cancellations but rather only future net bookings (as explained earlier), the overbooking method applied with CM2 and CM3 cannot be applied with CM1.

		CM2				CM3			
OBSCL		**0**	**0.5**	**1**	**1.5**	**0**	**0.5**	**1**	**1.5**
Ticket Revenue		0.12%	1.15%	2.08%	3.03%	0.10%	1.15%	2.07%	3.13%
DBs per 10K		0	10	71	190	0	10	71	187
Net Rev. (by DB cost)	**$150**	0.12%	1.10%	1.72%	2.08%	0.10%	1.10%	1.72%	2.19%
	$300	0.12%	1.06%	1.37%	1.13%	0.10%	1.06%	1.36%	1.26%
	$450	0.12%	1.01%	1.01%	0.18%	0.10%	1.01%	1.00%	0.32%
Load Factor		-0.49	1.05	1.95	2.02	-0.33	1.08	1.95	2.08
Yield		0.72%	-0.23%	-1.09%	-1.70%	0.50%	-0.26%	-1.11%	-1.63%

Figure 91: CM2/3 vs. CM1 result measures

Figure 91 summarizes the difference in the primary result measures of CM2 and CM3 versus CM1 under a medium cancellation rate scenario. In this scenario, 15% of all bookings eventually cancel before departure. This cancellation rate is similar to the cancellation rates reported by North American airlines. Cancellation forecasting could increase revenue gains by 0.12% even when no overbooking is applied (OBSCL=0). In this case, the load factors of CM2 and CM3 are lower compared to CM1 and, in turn, the yields are between 0.5% and 0.72% higher. When we increase the OBSCL, the overbooking gets more aggressive and this leads to more bookings being accepted in the system and higher ticket revenues. The ticket revenue gains of CM2/3 over CM1 range between 1.15% and 3.13%, depending on the OBSCL used. The load factors increase and the yields decrease the higher the OBSCL parameter chosen.

However, as in the real world, overbooking does have negative consequences if on departure day the number of passengers booked exceeds the number of seats on a flight leg. The passengers who did not board need to be compensated according to regulations and re-accommodated on later flights. The higher the number of denied boardings (DBs) is, the higher the costs are for the airline. The number of DBs when OBSCL=0.5 is approximately 10 (per 10,000 passengers) which is close to the airline industry's standard, however when OBSCL>0.5 the number of DBs is substantially higher. Taking in consideration the DBs and the (different) costs associated with them shows that cancellation forecasting and overbooking are still beneficial. In the most realistic scenario in which OBSCL is set to 0.5, the net revenue gain could range between 1.01% and 1.1%, depending on the DB costs. Overall, the results for CM2 and CM3 are very similar despite the difference approach as to which cancellation data is used.

CM4 uses the same approach for cancellation rate forecasting as CM3. However, while the overbooking methodology of CM3 only focuses on compensating for BIH cancellations, CM4 compensates for cancellations of BIH and BTC at the same time. Consequently, in order to prevent a double discount for cancellations of future bookings, CM4 uses gross BTC forecast (a forecast that does not take future cancellations into account) as an input to the seat allocation optimizer.

		CM4			
	OBSCL	**0**	**0.5**	**1**	**1.5**
Ticket Revenues		-0.34%	1.28%	2.74%	4.16%
DBs per 10K		0	12	101	273
Net Revenues (by DB cost)	**$150**	-0.34%	1.22%	2.23%	2.79%
	$300	-0.34%	1.15%	1.73%	1.43%
	$450	-0.34%	1.09%	1.23%	0.06%
Load Factor		-3.1	0.28	2.01	2.4
Yield		3.57%	0.79%	-0.96%	-2.17%

Figure 92: CM4 vs. CM1 result measures

Figure 92 summarizes the difference in the primary result measures of CM4 versus CM1 (medium demand and medium cancellation levels). CM4 performs worse than CM1 when overbooking is not applied due to the fact that the forecast input is gross BTC forecast which is higher than the forecast input in CM2 and CM3. In this case, the

RM system is very restrictive and rejects low fare bookings which translates to significantly lower load factor but higher yield. When overbooking is applied, the ticket revenue gains range between 1.28% and 4.16%. These revenue gains are higher than the revenue gains of CM2 or CM3. At the same time, the number of DBs with CM4 is higher, especially when OBSCL>0.5. This is due to the overbooking methodology of CM4, which overbooks more aggressively than CM3 as it not only compensates for BIH cancellation but for BTC cancellations as well. Despite the higher DBs, the net revenue gains of CM4 over CM1 are slightly higher than those of CM2 or CM3 and range between 1.09% and 1.22% (depending on the DB costs assumed by the airline) in the realistic scenario where OBSCL is set to 0.5. In comparison to CM2 and CM3, the magnitude of changes in the load factors and yields due to an increase of the OBSCL is greater.

This thesis also looked at the revenue gains of cancellation forecasting and overbooking under a high cancellation rate setting, as an attempt to address the different cancellation rates reported by different airlines. In the medium CXL rate scenario total cancellations summed up to approximately 15% of gross bookings, while in the high cancellation rate scenario the cancellations summed up to almost 30% of gross bookings. Figure 93 summarizes the difference in the primary result measures of CM3 and CM4 versus CM1.

		CM3			CM4		
OBSCL		**0**	**0.5**	**1**	**0**	**0.5**	**1**
Ticket Revenues		0.23%	3.48%	6.31%	-2.40%	3.69%	8.60%
DBs		0	5	133	0	20	307
Net Revenues (by DB cost)	**$150**	0.23%	3.45%	5.65%	-2.40%	3.59%	7.09%
	$300	0.23%	3.42%	5.00%	-2.40%	3.49%	5.57%
	$450	0.23%	3.40%	4.34%	-2.40%	3.39%	4.06%
Load Factor		-1.02	3.37	6.65	-8.68	0.72	7.24
Yield		1.52%	-0.69%	-3.11%	9.37%	2.61%	-3.49%

Figure 93: CM3/4 vs. CM1 results in a high CXL rates scenario

The trends in Figure 93 are similar to the trends in Figures 91 and 92. That is, the higher the OBSCL the higher the (ticket and net) revenue gains, DBs and load factors are, and lower the yields are. However, the main difference is the magnitude of the changes compared to CM1. In a high cancellation rate scenario, the revenue gains from cancellation forecasting are greater and range between 3.4% and 3.59% with either CM3 or CM4 in the realistic scenario where OBSCL is set to 0.5. These results show airlines with high cancellation rates could gain more from cancellation forecasting and overbooking than airlines with lower cancellation rates.

As in the real world, the result measures of an airline are not dependent only its own RM cancellation forecasting and overbooking strategy, but also the RM strategies of the competitors. In this thesis, we tested the changes in results after before and after the competitors move from using CM1 without overbooking (and thus not actually forecasting for cancellations) to CM2 with moderate overbooking (OBSCL=0.5). This transition led to revenue losses (for AL1) of 0.4% to 0.8% in ticket revenues and 0.65% to 0.73% in net revenues, depending on the OBSCL used. When the competition increases the availability of seats (especially in lower fare classes) due to overbooking, our RM system responds by opening more seats in the lower fare classes as well. The revenue loss is then a combination of two effects. First, a spiral down effect, which means passengers who would otherwise book in higher fare classes now book in lower fare classes instead. Second, a spill effect, which means passengers who would otherwise book with a certain airline, now book tickets with the competitors. The loss of passengers to the competitors, and in particular high fare class passengers, translated then to lower load factors and lower yields.

Changing the aggregation level of the cancellation rate estimation model from path (based on all possible origin/destination combinations in the network) to leg (based on all possible direct flights in the network) can lead to an increase of up to 0.25% in ticket revenues. The tests showed that the overbooking with leg cancellation rate estimates is slightly more aggressive. This is a direct consequence of the finding that path aggregated cancellation estimates tend to (slightly) underestimate cancellation rates while leg aggregated cancellation estimates tend to do the opposite. Since the overbooking methodology used in this test assumes a linear relationship between cancellations and overbooking, the higher the forecasted cancellation rate is, the more aggressive the overbooking is. This, in turn, results in higher number of DBs. The leg aggregated cancellation rate estimates have a fairly marginal net revenue gain over the path aggregated cancellation rates in the range of 0.03% and 0.13%, depending on the OBSCL and DB cost. When the DB costs are high and OBSCL>0.5, the net revenue impacts are negative.

The tests in this thesis showed that the deterministic overbooking mechanism used in all cancellation forecasting methods leads to high numbers of DBs compared with industry standard when OBSCL>0.5. In order to reduce the number of DBs, an "apex overbooking" mechanism was developed that caps the authorized capacity (AU) during the booking process. The apex overbooking mechanism does limit the number of DBs when OBSCL>0.5 to be around 20 (per 10,000). Even though this number is still more than double the industry standard, it is substantially smaller than the DB numbers based on the simplistic overbooking mechanism. However, the attempt to limit DBs also results in lower revenues as less bookings are accepted, especially in lower fare classes, compared with the previous overbooking mechanism. In turn, the results also showed lower load factors but higher yields.

Testing the performance of optimizers and forecasters while taking into account passenger cancellation behavior showed results that are similar to the results in previous work that excluded cancellations. The three different commonly used RM seat allocation optimizers that were compared are: Displacement Adjusted Virtual Nesting (DAVN), Probabilistic Bid Price (ProBP) and Unbucketed Dynamic Programming (UDP). The two different forecasters that were compared were Standard Forecasting (SF) and Hybrid Forecasting (HF). The hybrid forecasting was also used tested with the Fare Adjustment (HF/FA) technique which is targeted at preventing the spiral down effect. UDP has the highest ticket revenues regardless of the CXL forecasting method or the OBSCL used, while also having the lowest DBs compared to ProBP and DAVN. Consequently UDP has the highest net revenues. On the other hand, UDP has slightly lower net revenue gains over CM1 compared with DAVN and ProBP in all scenarios.

HF/FA had the highest ticket revenues and the lowest number of DBs compared with SF and HF. As result, the highest net revenues are gained with HF/FA. The RM system is more restrictive when HF or HF/FA is employed, hence the load factors are lower and yields are higher compared with SF.

The advancement in computational capabilities has pushed both academia and the industry to develop more sophisticated tools for cancellation forecasting purposes, using data on a disaggregate level. In the context of cancellation forecasting, the elements included in the airline's detailed Passenger Name Record (PNR) data can provide additional insights with regard to the cancellation behavior of airline passengers. Several examples for attributes that could be included in the analysis of cancellation behavior were described in this thesis. These attributes (in additional to other attributes) could be used in a logit model to calculate cancellation rates by leg or path, by class and time frame.

Despite the several approaches suggested in this thesis for cancellation forecasting and overbooking, not all possible approaches were covered. In particular, the overbooking model based on overbooking scaling (OBSCL) assumed cancellation rates are known with certainty. This model would be a reasonable solution if an airline is indifferent about the number of DBs and spoiled seats (seats left unoccupied at departure). In the real world, cancellation rates are not known with full certainty and hence future approaches should address this issue by developing probabilistic or risk-based overbooking models, as suggested in Belobaba (2015). Since different airlines have different cost estimates for DBs and spoiled seats, the optimal overbooking solutions will differ by airlines. In addition, this thesis did not directly address the issue of "no-shows" (booked passengers who fail to show up at the gate on time of departure) which is a more complex scenario for the airlines, as they have do not have the time to compensate for the potential loss of revenues as with cancellations occurring early in the booking process.

In conclusion, cancellation forecasting can help airlines increase their revenues. The benefits of cancellation forecasting are greater for airlines with higher cancellation rates (due to their less restrictive fare structure, for example). According to the simulation results, this is true also when overbooking is not applied. There could be different approaches for cancellation forecasting using different types of historical data. However, the results would not be significantly different if an airline chooses one method over the other. Airlines could increase their revenues even further if overbooking is applied. Airlines should not be too aggressive with overbooking, otherwise the benefits would be offset by the costs of compensating denied boardings. In addition, while the natural outcome of overbooking is an increase in load factors, increasing availability of seats (especially in low fare classes) could substantially decrease yields. Since both these metrics are constantly under scrutiny by airline managements, it is important to not to overbook too early in the booking process.

Bibliography

Belobaba, P. (1987). Air travel demand and airline seat inventory management. Cambridge, MA: Flight Transportation Laboratory, Massachusetts Institute of Technology.

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

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

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

Belobaba, P. (2002). 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.

Belobaba, P., & Hopperstad, C. (2004, June). Algorithms for revenue management in unrestricted fare markets. In Meeting of the INFORMS Section on Revenue Management, Massachusetts Institute of Technology.

Belobaba, P. (2010). Passenger Origin Destination Simulator (PODS): Summary of Processes and Models. Unpublished Report, Massachusetts Institute of Technology (September).

Belobaba, P., Odoni, A., & Barnhart, C. (2015). *The global airline industry*. John Wiley & Sons.

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

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

Chatterjee, H., (2001). Forecasting for cancellations. In: Presentation at AGIFORS 2001 Reservations and Yield Management Conference, Bangkok, Thailand.

Diwan, S. (2010). Performance of Dynamic Programming Methods in Airline Revenue Management. Master's thesis, Massachusetts Institute of Technology, Department of Civil and Environmental Engineering and Sloan School of Management.

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

Garrow, L. A., & Koppelman, F. S. (2004). Predicting air travelers' no-show and standby behavior using passenger and directional itinerary information. *Journal of Air Transport Management, 10*(6), 401-411.

Gosavi, A., Ozkaya, E., & Kahraman, A. F. (2007). Simulation optimization for revenue management of airlines with cancellations and overbooking. OR Spectrum, 29(1), 21-38.

Ignaccolo, M., & Inturri, G. (2000). OVERBOOKING IN AIR TRANSPORTATION. *Journal of Air Transportation World Wide, 5*(2-2000).

Iliescu, D. C., Garrow, L. A., & Parker, R. A. (2008). A hazard model of US airline passengers' refund and exchange behavior. *Transportation Research Part B: Methodological, 42*(3), 229-242.

Karaesmen, I., & van Ryzin, G. (2004). Overbooking with substitutable inventory classes. *Operations Research, 52*(1), 83-104.

Lautenbacher, C. J., & Stidham Jr, S. (1999). The underlying Markov decision process in the single-leg airline yield-management problem. *Transportation Science, 33*(2), 136-146.

Lawrence, R. D., Hong, S. J., & Cherrier, J. (2003, August). Passenger-based predictive modeling of airline no-show rates. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 397-406). ACM.

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

Lemke, C., & Gabrys, B. (2008). Forecasting and Forecast Combination in Airline Revenue Management Applications.

Lemke, C., Riedel, S., & Gabrys, B. (2009, March). Dynamic combination of forecasts generated by diversification procedures applied to forecasting of airline cancellations. In *Computational Intelligence for Financial Engineering, 2009. CIFEr'09. IEEE Symposium on* (pp. 85-91). IEEE.

Littlewood, K. (2005). Special issue papers: Forecasting and control of passenger bookings. *Journal of Revenue and Pricing Management, 4*(2), 111-123.

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

Morales, D. R., & Wang, J. (2010). Forecasting cancellation rates for services booking revenue management using data mining. *European Journal of Operational Research, 202*(2), 554-562.

Neuling, R., Riedel, S., & Kalka, K. U. (2004). New approaches to origin and destination and no-show forecasting: Excavating the passenger name records treasure. *Journal of Revenue and Pricing Management, 3*(1), 62-72.

Richter, H. (1982, October). 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.

Sierag, D. D., Koole, G. M., van der Mei, R. D., van der Rest, J. I., & Zwart, B. (2015). Revenue management under customer choice behaviour with cancellations and overbooking. *European Journal of Operational Research, 246*(1), 170-185.

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

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

Wang, K. (1982, October). 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.

Williamson, E. L. (1992). Airline network seat inventory control: Methodologies and revenue impacts. [Cambridge, Mass.: Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, Flight Transportation Laboratory.

Wysong, R. (1988, October). 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.