

Investigation of Competitive Impacts of Origin-Destination Control Using PODS

by

Alex Yen Hung Lee

S.B., Aeronautics and Astronautics (1996)

Massachusetts Institute of Technology

Submitted to the Department of Aeronautics and Astronautics
in Partial Fulfillment of the Requirements for the Degree of
Master of Science in Aeronautics and Astronautics

at the

Massachusetts Institute of Technology

June 1998

©1998 Massachusetts Institute of Technology
All rights reserved

Signature of Author

Department of Aeronautics and Astronautics
May 22, 1998

Certified by

Peter Paul Belobaba
Associate Professor of Aeronautics and Astronautics
Thesis Supervisor

Accepted by

Jaime Peraire
Chairman, Department Committee on Graduate Students

JUL 08 1998

LIBRARIES

AERO

Investigation of Competitive Impacts of Origin-Destination Control Using PODS
by
Alex Yen Hung Lee

Submitted to the Department of Aeronautics and Astronautics
on May 22, 1998 in Partial Fulfillment of the Requirements for the Degree of
Master of Science in Aeronautics and Astronautics

Abstract

The Passenger Origin / Destination Simulator (PODS) was used to investigate the competitive impacts of Origin-Destination control airline Revenue Management (RM) methods. Experiments performed included revenue performance of O-D control RM methods versus EMSRb fare class control, impacts of passenger choice assumptions, competing airlines with co-located hub vs. separate hubs, and different implementations of the most promising RM methods.

First, the revenue results from several O-D control RM methods are compared to the revenue result from EMSRb fare class control. Of the four O-D methods investigated, Displacement Adjusted Virtual Nesting (DAVN) performed consistently the best, followed closely by Heuristic Bid-Price (HBP). Greedy Virtual Nesting (GVN) was the third best performer, with decreasing performance as demand increased. Network Bid-Price (Netbid) performed the worst because of the “small network effect”, and due to its sensitivity to forecasting and detrunctation methods.

Next, the impacts of passenger choice assumptions were investigated. We found that with the inclusion of demand correlation and passenger choice in the simulations, DAVN and HBP performed better, while Netbid performed worse. The result for GVN was inconsistent. Under the new PODS maximum willingness-to-pay formulation with lower levels of sell-up, we saw lower passenger loads with more spill. We also noticed lower absolute revenue, as well as bigger revenue impact from O-D methods.

For competing airlines with co-located hub vs. separate hubs, most RM methods performed better under the co-located hub scenario tested because the airlines have to compete for both local and connecting passengers, making passenger selection more critical. As to O-D methods used by airlines competing head to head, under HBP and GVN, the co-located hub network leads to more sell-up, leading to better revenue performance.

Finally, different implementations of HBP and DAVN were investigated. For HBP, we found that the d-factor had a large effect on performance. For DAVN implementations, we noticed that different virtual class boundary methods had a large impact on performance, while frequent re-optimization might actually reduce performance.

Thesis Supervisor : Peter Paul Belobaba
Title : Associate Professor of Aeronautics and Astronautics

Acknowledgements

First, I would like to thank my thesis advisor, Professor Peter Belobaba, for his support and guidance for the past two years. He is undoubtedly the best advisor I could ever have imagined. I really appreciate his putting up with me and answering all the stupid questions I had. Again, thanks Peter.

I would like to thank my sponsors – Continental Airlines and KLM. Without the sponsorship of the two forward thinking airlines, this thesis would not have been possible.

Next, I would like to express my appreciation for all those who helped me with my PODS research. I would especially like to thank Dr. Craig Hopperstad of The Boeing Company for all the guidance on the PODS research, as well as detail explanations on all aspects of PODS.

Also, I would like to thank my friends and colleagues at MIT Flight Transportation Lab: Stephane Bratu, Aamer Charania, Michael Clarke, Daniel Skwarek, Julie Xie, and Jeff Zickus. Thanks for all the discussions about the airline industry, as well as the friendship during the past two years.

Furthermore, I would like to thank all my church friends for their endless care. I would especially like to thank Charles and Peter for their help on many things.

Finally, I would like to thank my family for all their love and support. I have to thank Mom for making sure I did my homework ever since I was a little kid. I have to thank Dad for providing us with a happy life. I have to thank Anne for being such a sweet sister. Finally, I especially have to thank Jane for helping me get through this hectic but memorable past week, as well as all her love and support throughout the years.

With all my gratitude, I dedicate this thesis to my family and friends.

Table of Contents

1	Introduction.....	13
1.1	Introduction to Revenue Management and Network Origin-Destination Control.....	13
1.1.1	Revenue Management.....	13
1.1.2	Network Origin-Destination Control.....	14
1.2	Objective of Thesis.....	15
1.3	Questions for Network Origin-Destination Control.....	15
1.3.1	The Relative Performance of Network Origin-Destination Control Revenue Management Methodology Compared to Leg-Based EMSRb Fare Class Control Methodology.....	15
1.3.2	Impacts of Passenger Choice Assumptions.....	16
1.3.3	Competing Airlines with Co-Located Hub vs. Separate Hubs....	16
1.3.4	Different Implementations of DAVN and HBP O-D Seat Inventory Control Methods.....	16
1.4	Structure of Thesis.....	17
2	Simulation Tool – The Passenger Origin / Destination Simulator (PODS)	19
2.1	Basic Structure of PODS.....	20
2.2	Passenger Choice Model.....	22
2.2.1	Generate Disutility Costs, Favorite Airline, and Decision Window.....	24
2.2.2	Generate Maximum Willingness-To-Pay (WTP)	24
2.2.3	Compute Total Perceived Cost.....	25
2.2.4	Allocation Process Amongst Available Path/Fare Classes.....	25
2.3	Revenue Management Seat Inventory Control.....	25
2.3.1	Expected Marginal Seat Revenue (EMSRb).....	27
2.3.2	Greedy Virtual Nesting (GVN).....	28
2.3.3	Displacement Adjusted Virtual Nesting (DAVN).....	31
2.3.4	Network Bid-price (Netbid).....	32
2.3.5	Heuristic Bid Price (HBP).....	33
2.4	Changes to PODS in Versions 6 and 7.....	35
2.4.1	Network Control – PODS6A.....	35
2.4.2	First Choice Only Choice (FCOC) – PODS6B.....	36
2.4.3	Disutility and Maximum Willingness-To-Pay – PODS6C.....	36
2.4.4	Forecast Estimation Categories – PODS6D.....	36
2.4.5	IBLV, Re-Optimization – PODS6E.....	37
2.4.6	Retro Crystal Ball – PODS6F.....	37
2.4.7	DAVN Implementations – PODS6I.....	37

	2.4.8	Different Maximum Willingness-To-Pay Formulation – PODS7A.....	38
2.5		Chapter Summary.....	39
3		Simulation Environment.....	41
3.1		System Level Inputs.....	41
	3.1.1	Top Level Settings.....	41
	3.1.2	Amount of Data Collected.....	43
	3.1.3	Demand Variability Parameters.....	44
	3.1.4	Disutility Cost Parameters.....	45
3.2		Airline Level Inputs.....	46
	3.2.1	Fare Class Structures.....	46
	3.2.2	Revenue Management System.....	48
3.3		Market / Network Level Inputs.....	52
	3.3.1	Network 1.....	52
	3.3.2	Network 2.....	54
	3.3.3	Network 3.....	55
	3.3.4	Parameters Common to the Three Networks.....	57
3.4		PODS Testing Sequences.....	58
	3.4.1	Sequence 0.....	58
	3.4.2	Sequence 2.....	59
	3.4.3	Sequence 3.....	59
3.5		Chapter Summary.....	60
4		Results and Discussion.....	61
4.1		The Relative Performance of Network Origin-Destination Control Revenue Management Methodology Compared to EMSRb Fare Class Control Methodology	62
	4.1.1	Motivation.....	62
	4.1.2	Simulation Setup.....	62
	4.1.3	Results and Discussions.....	63
	4.1.4	Summary.....	77
4.2		Impacts of Passenger Choice Assumptions.....	82
	4.2.1	Motivation.....	82
	4.2.2	Simulation Setup.....	82
	4.2.3	Sequence 0 and Sequence 2 vs. Sequence 3.....	83
	4.2.4	PODS7 Representation of Passenger Choice.....	87
4.3		Competing Airlines with Co-located Hub vs. Separate Hubs.....	96
	4.3.1	Motivation.....	96
	4.3.2	Simulation Setup.....	96
	4.3.3	Comparison Between Co-located Hubs and Separate Hubs.....	98
	4.3.4	Summary.....	103

4.4	Different Implementations of Different Revenue Management Seat Inventory Control Methods.....	105
4.4.1	Motivation.....	105
4.4.2	Heuristic Bidprice (HBP) Implementations.....	106
4.4.3	DAVN Implementations.....	107
5	Conclusion	117
5.1	Summary of PODS Investigation Results.....	117
5.2	Future Research Directions.....	121
	Bibliography	123
	Appendix	
A	PODS Input File Descriptions.....	125
B	A Sample PODS Input File.....	131
C	Passenger Loads and Class Distribution Under PODS6 vs. PODS7.....	133

List of Figures

2.1	Basic Schematic for PODS.....	20
2.2	PODS Passenger Choice Model.....	23
3.1	Booking Curves for Business and Leisure Passengers.....	42
3.2	Network 1 – Two airlines co-located their hubs, competing for both local and connecting traffic.....	52
3.3	Network 2 – Two airlines, each with their own hubs, competing only for O-D traffic.....	54
3.4	Network 3 – Two airlines with separate hubs, competing only for connecting traffic.....	55
4.1	The performance of O-D control RM methods under Network 1.....	78
4.2	The performance of O-D control RM methods under Network 2.....	79
4.3	The performance of O-D control RM methods under Network 3.....	80
4.4	% Revenue Changes Compared to EMSRb Under PODS6 vs. PODS7, DF = 0.8.....	90
4.5	% Revenue Changes Compared to EMSRb Under PODS6 vs. PODS7, DF = 1.0.....	92
4.6	% Revenue Changes Compared to EMSRb Under PODS6 vs. PODS7, DF = 1.2.....	94
4.7	Topology of Network 1.....	97
4.8	Topology of Network 2.....	97

List of Tables

2.1	Sample Network Used for RM Algorithm Example.....	26
2.2	Example Network Using the EMSRb Algorithm.....	27
2.3	Example Network Using the GVN Algorithm.....	29
2.4	Example Network Using the DAVN Algorithm.....	31
2.5	Example Network Using the Netbid Algorithm.....	32
2.6	Example Network Using the HBP Algorithm.....	34
2.7	Summary of Seat Allocation Algorithms.....	35
3.1	Results from the number of trials, number of samples per trial, and number of samples burned experiment.....	44
3.2	Fares, disutility costs, and perceived fares for business and leisure travelers in a 1000 mile market.....	48
3.3	Network 1 Leg Distances.....	53
3.4	Demand Summary for Network 1.....	53
3.5	Network 3 Leg Distances.....	55
3.6	Demand Summary for Network 3.....	56
3.7	Fare Chart for PODS Networks.....	57
3.8	Virtual Classes Ranges.....	57
4.1	Revenue Summary, $DF = 0.8$, Networks 1, 2, and 3.....	63
4.2	Revenue Summary, $DF = 1.0$, Networks 1, 2, and 3.....	65
4.3	Revenue Summary, $DF = 1.2$, Networks 1, 2, and 3.....	66
4.4	EMSRb vs. EMSRb Leg Passenger Load by Class Results.....	68
4.5	GVN vs. EMSRb Leg Passenger Load by Class Results.....	68
4.6	Netbid vs. EMSRb Leg Passenger Load by Class Results.....	69
4.7	DAVN vs. EMSRb Leg Passenger Load by Class Results.....	70
4.8	HBP vs. EMSRb Leg Passenger Load by Class Results.....	71
4.9	EMSRb vs. EMSRb Base Case Path Passenger Load by Class Results.....	73
4.10	GVN vs. EMSRb Base Case Path Passenger Load by Class Results.....	73
4.11	Netbid vs. EMSRb Base Case Path Passenger Load by Class Results.....	74
4.12	DAVN vs. EMSRb Base Case Path Passenger Load by Class Results.....	75
4.13	HBP vs. EMSRb Base Case Path Passenger Load by Class Results.....	76
4.14	Sequences 0 vs. 2 vs. 3, $DF = 0.8$	84
4.15	Sequences 0 vs. 2 vs. 3, $DF = 1.0$	85
4.16	Sequences 0 vs. 2 vs. 3, $DF = 1.2$	86
4.17	Demand Elasticity with Respect to Price for Business and Leisure Travelers...	89
4.18	PODS6 vs. PODS7 Revenue Results, $DF = 0.8$	90
4.19	PODS6 vs. PODS7 Revenue Results, O-D RM Methods Competing Head to Head, $DF = 0.8$	91
4.20	PODS6 vs. PODS7 Revenue Results, $DF = 1.0$	91
4.21	PODS6 vs. PODS7 Revenue Results, O-D RM Methods Competing head to Head, $DF = 1.0$	93
4.22	PODS6 vs. PODS7 Revenue Results, $DF = 1.2$	93
4.23	PODS6 vs. PODS7 Revenue Results, O-D RM methods Competing Head to Head, $DF = 1.2$	95

4.24	Network 1 vs. Network 2, DF = 0.8, Sequences 0 and 2.....	98
4.25	Network 1 vs. Network 2, DF = 0.8, Sequence 3.....	99
4.26	Network 1 vs. Network 2, DF = 1.0, Sequences 0 and 2.....	100
4.27	Network 1 vs. Network 2, DF = 1.0, Sequence 3.....	101
4.28	Network 1 vs. Network 2, DF = 1.2, Sequence 3.....	102
4.29	HBP with d-factors at 0.15, 0.40, and 0.25.....	106
4.30	DAVN(14) vs. DAVN under Network 1.....	109
4.31	DAVN Implementation Nomenclature.....	110
4.32	Virtual Class Upper Boundaries.....	111
4.33	Different Virtual Class Boundary Methods.....	112
4.34	DAVN(.8)(12) vs. DAVN(12) results.....	112
4.35	Different Re-optimization Schemes.....	113

Chapter 1

Introduction

1.1 Introduction to Revenue Management and Network Origin-Destination Control

1.1.1 Revenue Management

In today's competitive marketplace, both domestically and internationally, it is a lot more difficult for an airline to stay profitable. In order to increase profit for an airline, it has to reduce costs and increase revenue. To extract more revenue out of an airline's network, airlines started employing a strategy commonly known as revenue management. Since the cost of implementing revenue management is insignificant compared to its benefits, it becomes an easy way for an airline to increase its profitability. As a result, airlines all over the world have started implementing revenue management.

The term "Revenue Management" consists of two components – price differentiation and seat inventory control.¹ Price differentiation is the process in which different fare products with different restrictions are set in order to separate passengers into different categories. This is basically third degree price discrimination, and it separates passengers with different willingness-to-pay levels so that they will pay different fares. However, in the real world, price differentiation is dictated by the need to match what the competitors are doing with their fare products in order to stay competitive.

Since there is not much an airline can do about price differentiation, the second component of revenue management - seat inventory control – becomes much more important. It uses different algorithms to allocate the number of seats available to each

¹ Wei, Y. *Airline O-D Control Using Network Displacement Concepts*, May 1997, p. 11

fare product an airline offers on a given flight. A whole field of research has been focusing on finding the best seat inventory control algorithm for different situations in order to generate the highest revenue possible. It is in this area that this thesis concentrates its research on.

1.1.2 Network Origin-Destination Control

Up to this point, most airlines practice leg-based seat inventory control, meaning that the seat inventory allocation is determined for each leg independent of the other legs in the system. This is easier to do because airlines only have to keep data on each leg, instead of on origin-destination markets for the whole network. This whole seat inventory control process is on a much smaller scale when using leg-based control because one only has to optimize the seat allocation on one leg. However, the downside to this leg-based optimization is that local optimization does not lead to global optimization. Even if the seat allocation is optimal on each leg, over the whole network, it may not be optimal at all. This problem is further compounded by the development of hub-and-spoke networks, where most of the passengers travelling on the network make a connection at a hub city, thus traversing more than one leg.

The following example illustrates the reason leg-based seat allocation schemes may not be the optimal solution over a network. In this hypothetical case, our example airline has a flight that operates from city A to city C via city B. On this particular flight, the first leg from A to B has only one seat left; while the second leg from city B to C is very empty. There are two passengers, the first one is looking for a full fare Y class ticket to travel from A to B; while the second passenger is looking for a discount B class ticket to travel from A to C. The Y ticket from A to B costs \$500, while a discount B ticket, with the longer flight distance of A to C, costs \$700. In this case if this airline practices leg-based control, the last seat would have been sold to the Y passenger because on a per-leg basis, each full fare Y class passenger generates more revenue than each discount B class passenger. This would have netted the airline only \$500. On the other hand, if the airline practices network control, by comparing the revenue contribution of the two passengers,

this seat would have been sold to the discount passenger traveling from A to C, bringing in \$700 dollars. This example clearly illustrates why network origin-destination seat inventory control has an advantage over leg-based control.

1.2 Objective of Thesis

Airlines today realize the importance of network origin-destination control in their revenue management system, and there is a trend for them to rush into implementing it. However, it is not very well known what the exact effect of moving to network O-D control has on the different revenue management seat inventory control algorithms. Furthermore, more work has to be done to prove that network O-D control is actually beneficial to an airline currently exercising leg-based control. The objective of this thesis is to address the fact that the implications of network O-D control are still not yet very well known, and more experiments have to be done to prove that its benefit is real. Questions about network O-D control will be answered in this thesis, and the performance of different origin-destination revenue management seat inventory control methods will be compared. It is hoped that through the issues discussed in this thesis, the implications of revenue management origin-destination control can be better understood.

1.3 Questions for Network Origin-Destination Control

This following section presents a list of questions that are important to enhance the understanding of the impacts of network origin-destination control. This thesis attempts to explore the effects of network O-D control through raising questions and answering them using the Passenger Origin/Destination Simulator (PODS) simulation tool.

1.3.1 The Relative Performance of Network Origin-Destination Control Revenue Management Methodology Compared to Leg-Based EMSRb Fare Class Control Methodology

First, the obvious question – is network O-D control better than the current leg-based EMSRb fare class control scheme? Does an airline really benefit by moving from leg-

based EMSRb fare class control to network O-D control? What happens to airlines' revenue when one or more airlines move to network O-D control? Furthermore, is it a zero-sum game, meaning does the total revenue of all airlines in the market stay the same, or will it increase if airlines start exercising network O-D control?

1.3.2 Impacts of Passenger Choice Assumptions

Next set of questions is the issue of passenger choice. With today's hub-and-spoke networks, passengers have a lot more choice in their air travel. Traditional simulations indicate an improvement when exercising network origin-destination control; but are the results from these simulations realistic with real world passenger choice? This issue will be addressed by running experiments using the Passenger Origin/Destination Simulator (PODS) with ability to simulate and change different assumptions of passenger choice.

1.3.3 Competing Airlines with Co-Located Hub vs. Separate Hubs

In this section, the effects of competing airlines with the same hubs vs. competing airlines with different hubs will be discussed. The network in which two airlines operate the same hub city is reminiscent to the United/American situation at Chicago O'Hare; while airlines with different hubs is the situation for which the rest of the airlines in the world operate in. Are there any differences to the impacts of network O-D control vs. leg-based control? Which methods should an airline employ under both the "same hub" and the "different hub" situation?

1.3.4 Different Implementations of DAVN and HBP O-D Seat Inventory Control Methods

There are several different seat inventory control methods in use today, and each one has many different implementations. (Please refer to Section 2.3 for a more detailed explanation of various seat inventory control methods) An airline has to look at their own network and their data collection capabilities in order to make a decision on which implementation of which method is the best for them. In this section, the most promising O-D based seat inventory control methods, Displacement Adjusted Virtual Nesting

(DAVN) and Heuristic Bid-Price (HBP), will be discussed, and the impacts of different implementations will be addressed.

1.4 Structure of Thesis

Chapter 2 describes the simulation tool used for the research in this thesis – the Boeing Passenger Origin/Destination Simulator (PODS). First the difference between PODS and traditional revenue management simulators will be pointed out. Then, the main modules of PODS for which this thesis dealt with, namely the Passenger Choice Models and the Yield Management Controls module, are described in detail. Finally, advances in PODS Versions 6 and 7, which were made in order to perform the experiments detailed in this thesis, will be listed.

Chapter 3 lists the simulation setup for the experiments. First the general parameters for PODS simulation runs are presented, followed by a detailed explanation of the implications of each parameter relevant to the experiments. Finally the three networks used for the experiments, including the 4-city network, the decoupled hubs network, and the 6-city network, are introduced.

Chapter 4 presents the results of the PODS experiments, including answers to the questions raised in section 1.3. A detailed analysis of the results as well as their impacts will be presented. The phenomenon noted from the experiments will be explained and interesting issues will be discussed.

Finally, Chapter 5 concludes this thesis by summarizing the results of the joint Boeing/MIT research, and also summarizes the answers to the many questions raised about network origin-destination control. This thesis then concludes with a discussion of possible future research directions in order to further enhance the understanding of network O-D control.

Chapter 2

Simulation Tool – The Passenger Origin / Destination Simulator (PODS)

The computer tool used for the experiments in this thesis is called the Passenger Origin / Destination Simulator, or PODS. It is a computer simulation tool developed by The Boeing Company. This thesis is the third in a series of reports resulting from the research collaboration between Boeing and the MIT Flight Transportation Lab. A more detailed discussion of the structure and theory of the PODS simulation can be found in the two previous reports by Wilson (1995)¹ and Skwarek (1996).² In this thesis, a brief description of the PODS tool is provided, followed by a description of changes made to PODS after the Skwarek thesis.

The main difference between PODS and the more traditional revenue management simulations, such as that developed at MIT by Williamson³ (MITSIM), is that PODS is able to take into consideration the passenger choice models. In the traditional simulations, booking history is input into the simulation in order to test the performance of the underlying revenue management system. This is referred to as the “First-Choice-Only-Choice” (FCOC) because no path choice is allowed for each passenger. In the MITSIM, when a passenger wants to book in a particular class that is closed, that passenger gets spilled. On the other hand, in the real world, the passenger may decide to look at some alternatives, such as flying on a competing airline, or flying in a higher fare class. PODS breaks this FCOC barrier by allowing the analyst to specify the characteristics of passenger demand generated by the passenger choice model. This way,

¹ Wilson, J.L. *The Value of Revenue Management Innovation in a Competitive Airline Industry*, May 1995

² Skwarek, D.K. *Competitive Impacts of Yield Management System Components: Forecasting and Sell-up Models*, August 1996

³ Williamson, E.L. *Airline Network Seat Inventory Control: Methodologies and Revenue Impacts*, June 92

with a set of characteristics associated with each passenger generated, the simulation can then go through the path choices available and make a travel decision, just like a person would in real world. Furthermore, with its ability to generate a variety of networks and to incorporate revenue management algorithms, PODS is a versatile tool to investigate the impacts of many parameters that can be changed, as described in this chapter.

2.1 Basic Structure of PODS

The PODS simulation was created to simulate an environment in which two airlines compete with each other over a network defined by the user. The changes in PODS since the Skwarek thesis are related to the investigation in this thesis, which is to explore the competitive impacts of Revenue Management Origin-Destination control over a network. In this section, a simple description of the basic structure of PODS will be presented.

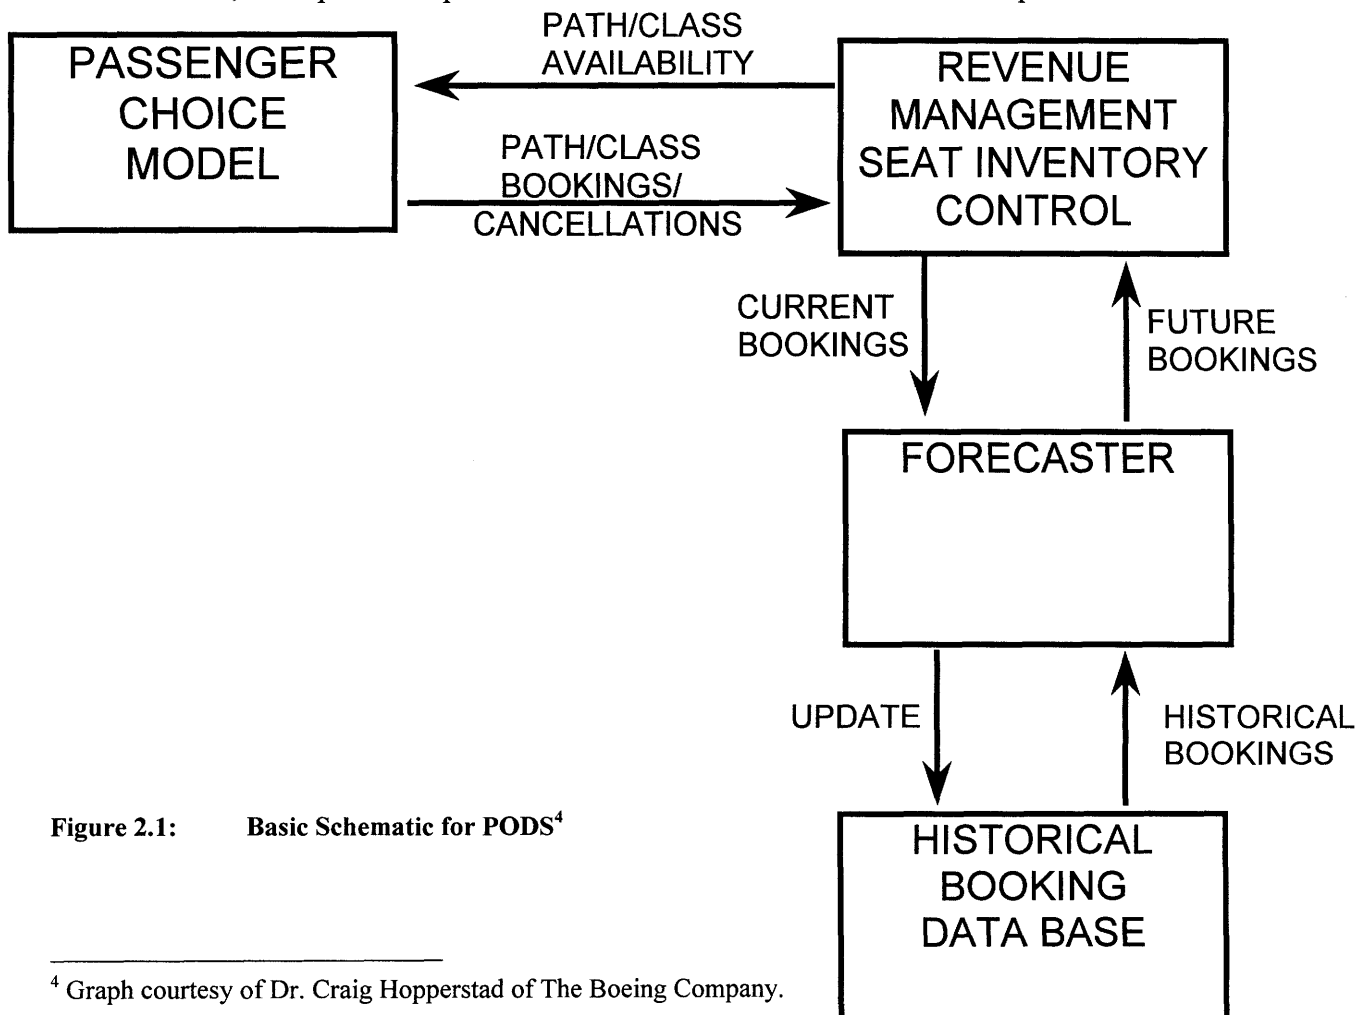


Figure 2.1: Basic Schematic for PODS⁴

⁴ Graph courtesy of Dr. Craig Hopperstad of The Boeing Company.

Please refer to Figure 2.1 for a schematic of the PODS simulation. The whole simulation starts in the Passenger Choice Model, which is the box on the top left. In this process, first, passengers for the PODS world are generated. They each have their own attributes and preferences. They can be either business passengers who are less sensitive to price, or leisure passengers who have a more elastic demand response. Next, according to how much the passengers are willing to pay, when they want to travel, and the disutility costs associated with each available fare, they make their choice as to which itinerary to travel on.

As passengers inquire about available fare classes before making their choice, the Revenue Management Seat Inventory Module is invoked. This is the box on the upper right in Figure 2.1. Each airline makes use of potentially different RM algorithms specified by the input parameters. Based on the information from the historical database and the forecaster, this module decides whether a particular fare is available for the itinerary requested. After the passenger decides which itinerary to travel on, he/she books the seat, then the Revenue Management Seat Inventory Module notes this fact and adjusts the seat availability accordingly.

This whole process is repeated for every single passenger generated by PODS, whose arrival rate is determined by the specified booking curve. The time line from the start of simulation to the departure date is divided into “time frames”, and at the end of each time frame, the Revenue Management Seat Inventory Control Module passes the current booking information to the forecaster. The forecaster then updates the historical booking database, and based on the historical booking, it generates a new forecast. Finally the Revenue Management Seat Inventory Control Module uses the new forecast to recalculate the availability of each fare class. This process is repeated for every single time frame, until the day of the flights. After the flights depart, PODS records the passenger load for each leg and OD market, as well as revenue generated. This result becomes one “sample”.

In order to increase the statistical significance of the results, multiple “samples” are generated as if there are multiple departure dates. The first sample uses the historical booking database input into PODS, while subsequent samples add in the generated data. Because the forecaster uses the user-input values to generate the first forecast instead of actual historical bookings, a number of samples are discarded, or “burned”. After a sufficient number of samples are generated, the results are averaged, and a “trial result” is produced. Multiple trials are then performed to come up with the final “case result”, which is then reported in this thesis. For every PODS experiment result that is reported, it is actually the “case result” described above, after going through the steps outlined above.

After describing the basic architecture of PODS, in the next two sections, major components of PODS relevant to the investigation of this thesis, including the Passenger Choice Model and the Revenue Management Optimizer, will be discussed in detail. This should provide a basis for the discussion about the experiment results presented in this thesis.

2.2 Passenger Choice Model

The ability to simulate passenger choice is what sets PODS apart from traditional Revenue Management simulations. Passengers make choices based on their willingness-to-pay, the price of the ticket, the flight time, and the availability of the seats. The Passenger Choice Model allows PODS to generate each passenger with a set of characteristics associated with that passenger, and let him/her go through the booking process to determine whether the passenger gets a seat or gets spilled. There are four steps involved in simulating the choice of a passenger. Please refer to Figure 2.2 for the schematic of the Passenger Choice Model.

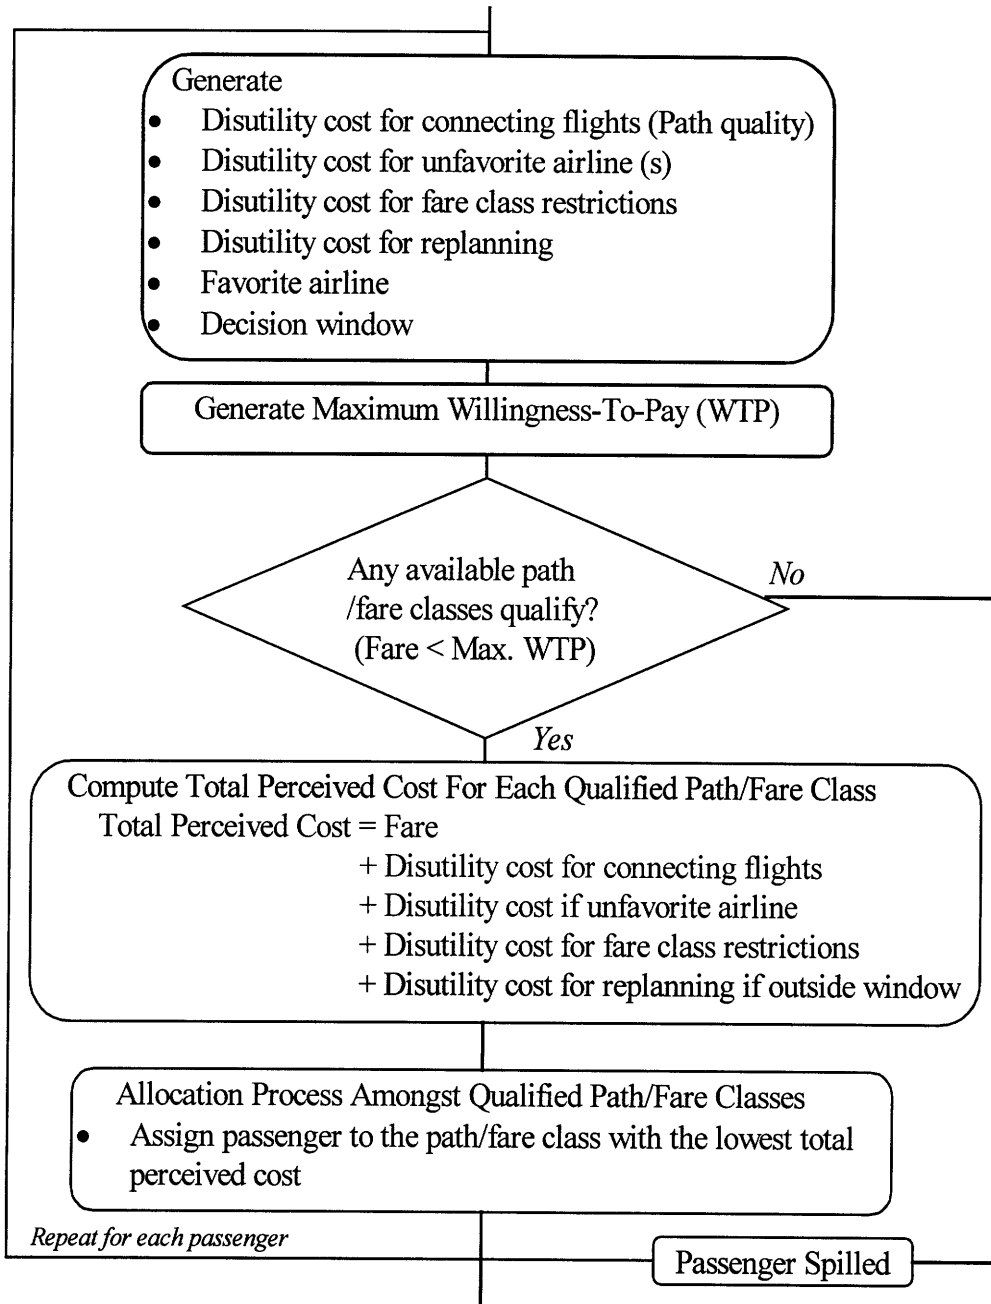


Figure 2.2: PODS Passenger Choice Model⁵

⁵ Special thanks to Dr. Hopperstad for comments and corrections on Figure 2.2

2.2.1 Generate Disutility Costs, Favorite Airline, and Decision Window

In the top block of Figure 2.2, first disutility costs, favorite airline, and decision window are generated for each passenger. There are four major categories of disutilities in PODS. They are disutilities for connecting flights (path quality), disutilities for unfavorable airlines, disutilities for different fare class restrictions, and disutilities for replanning (i.e. having to select an alternative decision window). Each disutility is modeled as a cost to each passenger, and later added onto the fare when the passenger makes his choice. Additionally, business passengers and leisure passengers have different disutility costs thus reflecting the difference in how the two passenger types value the disutilities.

As for the favorite airline and decision window, PODS assigns them based on the input to PODS for each passenger. A disutility cost is then assigned to the unfavorable airline as well as paths that are outside of the passenger's decision window.

2.2.2 Generate Maximum Willingness-To-Pay (WTP)

Next, a maximum willingness-to-pay value is generated. In the input file, the user specifies a willingness-to-pay factor for each type of passenger. At a fare of $[(\text{market base fare}) * (\text{willingness-to-pay factor})]$, only 50% of the passengers of this type are willing to travel. This allows PODS to construct a demand curve, which in turn generates the passenger demand for each trial. The shape of the demand curve is different in the different versions of PODS, and this change will be discussed in Section 2.4.8. This allows the experimenter to check for different passenger choice model assumptions, and see its impact on the revenue outcome and passenger mix of a market.

With the maximum willingness-to-pay value generated, all the paths and fare classes available are checked. If the dollar fare all the available path/fare classes for the passenger are higher than his/her maximum WTP value, then this passenger does not travel, and is "spilled". If there are available path/fares within his/her maximum WTP

value, then the simulation moves onto the next section to calculate the total perceived cost for those qualified path/fare classes.

2.2.3 Compute Total Perceived Cost

For each of the fares that qualify the maximum WTP requirement, its total perceived cost is calculated by the following formula:

$$\begin{aligned} \text{Total Perceived Cost} = & \text{Fare} \\ & + \text{Disutility cost for connecting flights} \\ & + \text{Disutility cost for unfavorite airline} \\ & + \text{Disutility cost for fare class restrictions} \\ & + \text{Disutility cost for replanning if outside decision window} \end{aligned}$$

This total perceived cost is then used to decide which of the available path/fare class the passenger chooses.

2.2.4 Allocation Process Amongst Available Path/Fare Classes

Finally, the results from the first three steps get passed onto the allocation process, which assigns passengers to available seats. For each passenger, the total perceived costs for all qualified paths are compared, and the path/fare class with the lowest perceived cost is chosen and the passenger reserves a seat on that path/fare class. By applying the same four steps to all passengers, PODS can then aggregate the loads and revenues for each flight and create an output.

In this section, a description of the passenger choice model was provided. A step-by-step explanation of how PODS generates the passenger choice was presented. The Passenger Choice Model is what sets PODS apart from traditional simulations.

2.3 Revenue Management Seat Inventory Control

There are five different major seat inventory control algorithms in PODS. They are the Expected Marginal Seat Revenue method (EMSRb), the Greedy Virtual Nesting method

(GVN), the Displacement Adjusted Virtual Nesting method (DAVN), the Network Bid-Price method (Netbid), and finally the Heuristic Bid-Price method (HBP). Detailed description of each algorithm can be found in Wei (1997)⁶. In this thesis, only a simple description of each method as it is implemented in PODS is provided. Examples will be used to illustrate the difference between the algorithms. All five methods are included in the experiments; therefore, it is important to understand the differences between each method.

A simple 2-leg Network is used as an example to illustrate how each of the algorithms work.

	Published Fare Class	Advanced Purchase Restrictions	Published Fares
Local	Y	Unrestricted	500
	B	7-day adv.	300
	Q	14-day adv.	150
Connection	Y	Unrestricted	800
	B	7-day adv.	400
	Q	14-day adv.	200

Table 2.1: Sample Network Used for RM Algorithm Example

Each of the 2 legs have local as well as connecting passengers, and each O-D market has 3 published fares available with different advance-purchase requirements. The published fares for both local markets are \$500 for the unrestricted Y class, \$300 for the 7-day advance purchase B class, and \$150 for the 14-day advance purchase Q class. For the connecting market, the fares are \$800, \$400, and \$200, respectively. (See Table 2.1) This same simple network will be used for the illustration of all five Revenue Management algorithms below.

⁶ Wei, Y. *Airline O-D Control Using Network Displacement Concepts*, May 1997, Chapter 2

2.3.1 Expected Marginal Seat Revenue (EMSRb)

The base case algorithm is the EMSRb algorithm. Normally when we refer to “EMSRb”, we are talking about a seat inventory control algorithm developed by Belobaba⁷. However, in the PODS world, EMSRb refers to a leg-based fare class control RM algorithm that uses EMSRb to optimize the seat inventory availability control for each leg.

This algorithm is also the most commonly used RM algorithm today. Since it is a leg-based fare class control algorithm rather than an Origin-Destination control algorithm, this makes it a good base case for comparing how well the O-D algorithms perform.

	Published Fare Class	Advanced Purchase Restrictions	Published Fares	Booking Classes	Data Collection/Forecast	Optimization /Availability Control
Local	Y	Unrestricted	500	Y	Y	Y - 100
	B	7-day adv.	300			
	Q	14-day adv.	150			
Connection	Y	Unrestricted	800	B	B	B - 60
	B	7-day adv.	400			
	Q	14-day adv.	200			

Table 2.2 – Example Network Using the EMSRb Algorithm

Because EMSRb is a leg-based algorithm, booking classes, data collection and forecasting, as well as optimization and availability control are all on a leg basis. With the example network, what the travel agents or customers see are the published fare classes – Y, B, and Q, with their associated advance-purchase restrictions as well as the published fares. Behind the scenes in the airline’s computer system, for each leg in this O-D market, Y, B, and Q classes bookings are created. Data collection and forecasting are done for the Y, B, and Q classes for each of the two legs separately, and booking limits are calculated using the now-standard EMSRb seat inventory allocation algorithm.⁸ In our example, the EMSRb booking limit for Y, B, and Q classes may be 100, 60, and 50, respectively, for leg 1. For a local passenger traveling on leg 1, if he wants to book a

⁷ Belobaba, P.P. *Air Travel Demand and Seat Inventory Management*, 1987

B class ticket, he will see that there are seats available to him (60 seats available), and he will pay \$300 to travel. If he books that B class ticket, one seat is deducted from the available seats in B and Q classes for leg 1. For a connection passenger traveling on both legs 1 and 2, if he wants to book a B class ticket, first B class has to be available on both legs 1 and 2. In the case when B class is available on leg 1 but not on leg 2, the passenger has to book in Y class and pay the \$800 fare. After he purchases the \$800 ticket, one seat will then be deducted from the available seats on both legs 1 and 2.

Since EMSRb does not distinguish between local and connecting passengers who have different revenue potential, the following scenario illustrates how an O-D based algorithm may be of an advantage in a network environment. For example if there is only one seat left on leg 1, but leg 2 is wide open, the airline is obviously better off taking a connecting Y passenger paying \$800 instead of a local Y passenger paying \$500. However, since the revenue management system does not distinguish between connecting and local Y passengers, if a local passenger shows up next, the airline will accept his/her booking and receive \$500 in revenue. On the other hand, if O-D algorithms are utilized, the reservation system can then realize that the connecting passenger will bring in \$300 more in revenue over the whole network, and reject that local passenger.

In summary, EMSRb as we refer to in the PODS world, is a leg-based booking class RM approach. It does its data collection, forecasting, optimization, and availability control, all on a per-leg basis. This can lead to lost revenue potential in a network environment. It uses the EMSRb algorithm for seat inventory optimization, thus the name.

2.3.2 Greedy Virtual Nesting (GVN)

Greedy Virtual Nesting is also known as VEMSRb, or Virtual-EMSRb. As its name implies, instead of booking classes being actual published classes, under VEMSRb the booking classes are now “virtual” classes. The name “greedy” comes from the fact that under Greedy Virtual Nesting, airlines always favor the highest fare passengers. Greedy

⁸ See Appendix 1 of Wei (1997) for a description of the calculation of EMSR curve

Virtual Nesting is the most basic O-D based algorithm. It still performs forecasting and seat inventory control on a per-leg basis.

	Published Fare Class	Advanced Purchase Restrictions	Published Fares	Booking Classes	Data Collection/ Forecast	Optimization /Availability Control
Local	Y	Unrestricted	500	↗ Y1	Y1	Y1 - 100
	B	7-day adv.	300	↘ Y2	Y2	Y2 - 90
	Q	14-day adv.	150	↘ Y3	Y3	Y3 - 80
Connection	Y	Unrestricted	800	↘ Y4	Y4	Y4 - 65
	B	7-day adv.	400	↘ Y5	Y5	Y5 - 40
	Q	14-day adv.	200	↘ Y6	Y6	Y6 - 15

Table 2.3: Example Network Using the GVN Algorithm

Table 2.3 shows our example network under the GVN algorithm. Its published fare classes and fares, as well as advance purchase restrictions, are all identical to the EMSRb case. Therefore, although the underlying RM algorithm is different, there is no difference to the customers or travel agents. What is different is behind the scenes in the airline's computer system. Instead of booking passengers in Y, B, and Q classes as published, the airline creates virtual booking classes for each leg. The PODS analyst specifies the number of virtual classes for each airline, and in our example network there are 6 virtual classes, labeled Y1 to Y6.

As for virtual class boundaries, the range of published fares in the network is first surveyed. This fare includes both local fares as well as connection fares. Then, the published fares are grouped into virtual booking classes, and boundaries established between the virtual booking classes. This set of boundaries is then employed on all legs throughout the network. As for data collection, forecast, optimization, and availability control, it is all done on a per-virtual class, per-leg basis. Seat inventory control booking limit optimization is again done using the EMSRb algorithm. Note that since different fare classes with different expiration dates may be grouped into the same virtual class, under the PODS implementation, the expiration of that particular virtual class is then the latest expiration date of all fare classes assigned to it.

Since the GVN algorithm is transparent to the customers, to implement GVN the airline has to have CRS seamless availability. The airline's internal computer has to determine in real time which virtual class a particular requested fare should be in, and in real time decide whether to accept that booking or not. Since this fare class to virtual class mapping is internal to the airline's computer system and transparent to the computer reservation system, they have to have seamless availability in order for GVN to work.

In our network example, each ODF is mapped into a different "virtual" fare class in the airline's computer reservation system. Connecting Y class with a \$800 fare is in virtual class Y1, while local Y class with a \$500 fare is in virtual class Y2. Again if leg 1 has only one seat left but leg 2 is wide open, under GVN that seat on leg 1 will be available to virtual class Y1. This ensures that the remaining seat is only available to the connecting Y passenger, who has to pay \$800 to travel.

This GVN algorithm has its drawbacks as well. If both legs 1 and 2 are very full, the airline is better off selling the seats to 2 local passengers instead of 1 connecting passenger. Under GVN, that seat will be sold to a connecting passenger because the connecting fare is in a higher virtual class, and this will bring in \$800 in revenue for the airline. However, if the airline could sell the seats to a local Y passenger on leg 1, and to another local Y passenger on leg 2, they will receive \$1000 in revenue. This is the case when the GVN algorithm is being too "greedy", and thus generating less revenue than our leg-based EMSRb algorithm.

In summary, Greedy Virtual Nesting is again a leg-based algorithm. Under GVN, there are virtual classes which published fare classes are mapped into. Data collection, forecasting, optimization, and availability control are all done on a per-leg, per virtual class basis. GVN always gives preference to connecting passengers over local passengers, and this may lead to an increase of revenue or a decrease of revenue over EMSRb, depending on how full the flights are. In order to implement GVN, the airline has to have CRS seamless availability because the availability of seats is decided in real time on a per request basis. The advantage of GVN is that it is transparent to the travel

agent, customers, as well as competitors. Its disadvantages are that it requires all new forecasts because of the new virtual classes, and the computer equipment may require an expensive upgrade in order to handle CRS seamless availability.

2.3.3 Displacement Adjusted Virtual Nesting (DAVN)

DAVN stands for Displacement Adjusted Virtual Nesting. It uses a deterministic network linear program (LP) to generate shadow prices for each leg for each departure, using Origin-Destination Fare class (ODF) forecast. This shadow price is then used to calculate a pseudo-fare for each leg, as follows:

Pseudo-fare for leg 1 = total fare – shadow price for leg 2

Pseudo-fare for leg 2 = total fare – shadow price for leg 1

Finally for each virtual class, Pseudo-fares are used instead of actual fares in EMSRb booking limit calculations.

	Published Fare Class	Advanced Purchase Restrictions	Published Fares	Pseudo Fares	Booking Classes	Data Collection/Forecast	Optimization	Availability Control
Local	Y	Unrestricted	500			by ODF	Calculate Shadow Price for each leg to generate Pseudo Fares	100 85 65 55 30 15
	B	7-day adv.	300		Y1	Y1		
	Q	14-day adv.	150		Y2	Y2		
Connection	Y	Unrestricted	800	650	Y3	Y3		
	B	7-day adv.	400	250	Y4	Y4		
	Q	14-day adv.	200	50	Y5	Y5		
					Y6	Y6		

Table 2.4: Example Network Using the DAVN Algorithm

Table 2.4 shows our example network under the DAVN algorithm. First, a deterministic LP is run to generate shadow prices for each leg. In our example, the shadow price for the 2nd leg is found to be \$150. Next, pseudo-fares are calculated by subtracting the shadow price from the total fare for multi-leg paths. Note that for multi-legged paths, displacement cost for each leg is set to the shadow price of the other leg. In our example network, Y, B, and Q pseudo fares are then \$650, \$250, and \$50, respectively, for the connection passengers. The local total fares, plus the connection pseudo fares, are then assigned to the virtual booking classes according to the virtual class boundaries. The

same boundaries are used for all legs for that airline in the network. Data collection and forecasting are on an ODF basis. Virtual class forecasts for each leg are then constructed from ODF forecasts by adding means and variances. Finally, the EMSRb algorithm is used to calculate booking limits for each virtual class on each leg.

In summary, an airline using DAVN performs their forecast and data collection by O-D path, and their seat inventory control by leg. Shadow prices for each leg are calculated using a deterministic LP, and pseudo-fares for each ODF are generated by subtracting the shadow prices of the other leg from the fare. This pseudo-fare is then assigned to virtual classes, and their booking limits calculated using EMSRb.

2.3.4 Network Bid-price (Netbid)

Network Bid-Price uses deterministic LP to generate bid prices for each leg for each departure, using ODF forecasts. Bookings are accepted when the total itinerary fare is higher than the sum of the shadow prices of the legs traversed.

	Published Fare Class	Advanced Purchase Restrictions	Published Fares	Booking Classes	Data Collection/Forecast	Optimization	Availability Control
Local	Y	Unrestricted	500	none	by ODF	Run LP for each departure data	accept booking if FARE > shadow price 1 + shadow price 2
	B	7-day adv.	300				
	Q	14-day adv.	150				
Connection	Y	Unrestricted	800	none	by ODF	Run LP for each departure data	accept booking if FARE > shadow price 1 + shadow price 2
	B	7-day adv.	400				
	Q	14-day adv.	200				

Table 2.5: Example Network Using the Netbid Algorithm

Table 2.5 shows our example under the Network Bid-Price algorithm. It performs data collection and forecast by ODF path/fare class. Unlike previous methods, Netbid does not use booking classes. Instead, the mean forecast path/fare class demand is used as deterministic path/fare class input to calculate the network solution to leg bid-price. All O-D fares greater than or equal to the sum of leg bid-prices to be traversed (shadow prices) are then flagged as available, and stay available until another re-optimization process is performed to update the bid prices.

In our example, at the beginning of the booking process for each departure, a deterministic LP is first run to determine the shadow prices, or leg bid-prices, for each leg in the network. Until another LP is run to update the shadow prices, any local fares greater than the shadow price for that leg, or any connection fares greater than the sum of the shadow prices of the legs traversed, are all accepted. Without updating the shadow prices, any fares greater than the sum of shadow prices will be accepted. Without re-optimization of the network, too many seats may be taken up by passengers who are paying a lower fare. Therefore, it is crucial for Netbid to update the shadow prices frequently. The Netbid implementation under PODS has a mechanism to trigger how often the re-optimization is performed. Please see Section 3.2.2 for how this re-optimization is done in PODS. Furthermore, since the shadow prices come from the deterministic ODF forecasts, different forecast and detrunctation methods can affect the shadow prices greatly, thus effecting the performance for Netbid. For the effect of forecasting and detrunctation on Netbid, please see Zickus (1998).

In summary, Network Bid-Price calculates shadow prices for each leg using deterministic LP. All bookings for fares higher than the sum of the shadow prices of the legs traversed are accepted. As the booking process goes on, shadow prices will then be updated to prevent too many low fare bookings being accepted. Finally, since forecasts can effect the shadow prices, different forecasting and detrunctation methods have a big effect on the performance of Netbid.

2.3.5 Heuristic Bid Price (HBP)

Heuristic Bid Price is also called the Greedy Virtual Nesting with EMSR Heuristic Bid Price. As its name suggests, it is very similar to Greedy Virtual Nesting. The only difference between GVN and HBP is in optimization and availability control. The main advantage of HBP is that airlines can stay with the existing leg-based data collection and forecasting system while still having some benefit of a network-based revenue management algorithm.

	Published Fare Class	Advanced Purchase Restrictions	Published Fares	Booking Classes	Data Collection/ Forecast	Optimization	Availability Control
Local	Y	Unrestricted	500	Y1	Y1	local booking limits	Local: EMSR class limit
	B	7-day adv.	300	Y2	Y2		
	Q	14-day adv.	150	Y3	Y3		
Connection	Y	Unrestricted	800	Y4	Y4	EMSRb value for each leg is used	CNX: accept booking if Fare > E1 + dE2 and Fare > E2 + dE1
	B	7-day adv.	400	Y5	Y5		
	Q	14-day adv.	200	Y6	Y6		
						Leg 1 = E1 Leg 2 = E2	

Table 2.6: Example Network Using the HBP Algorithm

Table 2.6 shows the example network using the HBP algorithm. HBP has the same six virtual booking classes as GVN. It also does its data collection and forecast on a per-virtual class, per-leg basis. For local paths, HBP uses EMSRb booking limits; for connection paths, there is no booking limit, and instead the expected marginal seat revenue (EMSR) value for the remaining booking capacity for each leg is used to determine whether a booking is accepted or not. For connections, a booking is accepted if the following formulas are satisfied:

$$\text{FARE} > \text{EMSR}_{\text{leg1}} + \text{d-factor} * \text{EMSR}_{\text{leg2}}$$

and

$$\text{FARE} > \text{EMSR}_{\text{leg2}} + \text{d-factor} * \text{EMSR}_{\text{leg1}}$$

The d-factor is a factor related to the proportion of local/connecting passenger mix. An experiment of different d-factor values has been done, and the results are presented in Section 4.5.2. All paths/fare classes with a fare satisfying the above conditions are then flagged as available, and remain available until the re-optimization process is run to update the EMSR values.

In summary, an airline using HBP can use the existing leg-based database and forecasting system without the need to change to an ODF database. The Expected Marginal Seat Revenue value for each leg is then used to determine whether a booking is accepted.

Of the above five methods, only the EMSRb is a pure leg-based method. GVN, DAVN, Netbid, and HBP are all O-D based methods. They either use a path-based data collection and forecast, or a path-based seat inventory control algorithm. Here is a summary of different Revenue Management algorithms:

RM algorithm	Data Collection and Forecast	Seat Inventory Control
EMSRb	Leg-based	Leg-based
GVN	Leg-based	Leg-based
DAVN	Path-based	Leg-based
Netbid	Path-based	Path based
HBP	Leg-based	Path-based

Table 2.7: Summary of Seat Allocation Algorithms

In this section, different revenue management seat inventory control algorithms, as well as the assumptions associated with them, were discussed. All five algorithms were used in PODS experiments for this thesis. In the following section, changes to PODS relating to the experiments in this thesis will be discussed.

2.4 Changes to PODS in Versions 6 and 7

The first five versions of PODS were only capable of simulating a point to point, single leg market. Version 6 adds the ability to run competing airlines over a network scenario, while different subversions of Version 6 add different capabilities as the research moves onto different issues outlines in section 1.3. Finally in Version 7 we saw a modification to the passenger choice model which gives the capability to simulate “pricing games” or “demand stimulation” through price. This section gives a thorough description of the evolution of PODS from version 6 to 7.

2.4.1 Network Control - PODS6A

PODS6A, the first version of PODS6, added the capability to simulate hub and spoke networks. This was accomplished by increasing all the dimensions in PODS, as well as some input and algorithm changes in order to accommodate the network scenarios.

2.4.2 First Choice Only Choice (FCOC) - PODS6B

In order to validate PODS results with the MITSIM, capabilities were built in to be able to turn off “passenger choice” and to reduce demand variance. With no passenger choice and low variance, PODS can approximate the traditional simulations. This reduced passenger choice is referred to First-Choice-Only-Choice (FCOC). By running simulations with different levels of passenger choice and demand variance, the PODS results could be compared to the MITSIM results to get a better understanding of the effect of PODS variance of demand as well as full path/fare class choice. The results and discussions of experiments regarding capabilities added in this version, such as FCOC, can be found in Section 4.2.

2.4.3 Disutility and Maximum Willingness-To-Pay - PODS6C

In previous versions of PODS where a network consists of only one city pair, the fare price was within a small range. This allowed the disutility values to be a fixed constant for each type of restrictions. However by moving onto the connecting network, the available fare vary a great deal, depending on the length of the connection. In this case, due to the large difference in fares, the same restriction could be worth a large portion, or an insignificant amount, of a ticket fare. In order to correct this problem, the disutility cost was changed from a constant value to a factor of the base fare, allowing the ratio of fare to disutility to stay the same throughout the network.

2.4.4 Forecast Estimation Categories - PODS6D

In order to investigate the effect of forecasting on a network basis vs. on a per-leg basis, capabilities were introduced in 6D to label each leg/path separately for forecasting purposes. The leg/paths with the same label will be forecasted together, thus allowing forecasting based on a single leg/path, on the whole network, or on a group of leg/paths. The results and discussions of experiments regarding capabilities added in this version can be found in Section 4.3.

2.4.5 IBLV, Re-Optimization - PODS6E

At this time the research was concentrated on the network bid price (Netbid) and heuristic bid price (HBP) algorithms. This functional improvements made in this version were all made for the above two methods. First, the IBLV parameter was added, to tell the revenue management system whether or not the forecast is “believed”. If it is, the forecast is decremented by the number of bookings, until the next re-optimization. Another parameter added was NREOP, which determines the frequency of re-optimization. Finally, the availability processor was introduced to Netbid and HBP in order to improve their performance. Without the availability processor, an airline could be flooded with bookings and thus giving too many seats to low fare passengers. With availability processor, this tendency could be controlled.

2.4.6 Retro Crystal Ball - PODS6F

This version introduced the “Retro crystal ball”. This allowed the PODS simulation to run almost exactly like a traditional revenue management simulator, without passenger choice and variation in demand. Another element introduced in this version was the ability to switch off the displacement cost for the DAVN algorithm. This allowed the effect of displacement cost to be realized by comparing the differences between DAVN with and without displacement cost.

2.4.7 DAVN Implementations - PODS6I

There was a lot of interest in the DAVN algorithm, which performed very well in PODS simulations. Two input parameters were added to enable different implementations of DAVN. The first parameter, methdsp, introduced the ability to implement different virtual class boundary settings. One setting allowed the user to input virtual class boundaries, while another setting allowed the boundaries to be set by equalizing the demand on each leg. The second parameter, methrvc, enabled the control of whether the virtual class boundaries and displacement costs were optimized once at the beginning of

each trial, or once at the beginning of each time frame. Please see Section 4.5.3 for the results and discussions of different DAVN implementations.

2.4.8 Different Maximum Willingness-To-Pay Formulation - PODS7A

In the original formulation of the passenger choice model, the demand was fixed at the level input by the user, regardless of the fares from the competing airlines. In order to simulate the demand stimulation effect, the user had to manually increase the demand as fare is dropped. PODS7 incorporated mechanisms to automatically increase demand as a result of lowering the fares.

Here is the new formulation for our maximum willingness to pay in our passenger choice model.

$$P(\text{xpays}_{m,t} \geq f) = \min \left[1., e^{-\left[\frac{.6931(f - \text{fb}_{m,t})}{(\text{em}_{m,t} - 1)\text{fb}_{m,t}} \right]} \right]$$

where:

- xpay_{m,t}** = max willing-to-pay, market m, pax type t
- fb_{m,t}** = input base fare, market m, pax type t
- f** = fare
- em_{m,t}** = input multiple of base fare (emult) at which the probability a random pax of type t in market m will travel = 0.5

This elasticity style formulation allows price excursions without the user having to change the demand manually in order to simulate the demand stimulation by a lower fare.

This new formulation also allows us to derive our elasticity of demand, which is:

$$(\varepsilon_{m,t})_{D,P} = -\left(\frac{f}{fb_{m,t}}\right) \times \left(\frac{0.6391}{em_{m,t} - 1}\right)$$

where:

$(\varepsilon_{m,t})_{D,P}$ = elasticity of demand with respect to price, for market m, pax type t.

The results from the investigation of this new passenger choice model will be presented in Section 4.2.4. The old formulation represented the maximum willingness-to-pay with a slope and intercept, and its levels were fixed regardless of fare level. This new version has an exponential curve, while the maximum willingness-to-pay curves moves up and down with changes in fares. The biggest impact of this exponential curve is that there is less sell-up potential because of its shape. This lower level of sell-up leads to a different fare class distribution. It also results in slightly different impact amongst different revenue management algorithms. The non-linearity of the maximum willingness-to-pay curve, together with allowing fare changes to stimulate passenger demand, this new formulation should be a better representation of the real world.

In this section, a detailed account of the evolution of PODS was presented. From Version 6A where networking capability was introduced, to Version 7A where a modified passenger choice model allowed the simulation of elasticity of demand, every change to the PODS software was chronologically listed.

2.5 Chapter Summary

In this Chapter, the simulation tool used for this Thesis, the Passenger Origin/Destination Simulator (PODS), was introduced. The topics discussed included the basic structure of PODS, a detailed description of the passenger choice model, followed by a list of the assumptions made for each revenue management seat inventory control algorithms, and

finally a chronology of PODS from Version 6 to Version 7. In the next Chapter, the input parameters to PODS used in this thesis will be discussed.

Chapter 3

Simulation Environment

In this chapter, the simulation environment for which all the experiments were performed is discussed. PODS divides its inputs into three sections – the system level inputs, the airline level inputs, and the market level inputs. (Please see Appendix A for a PODS input file description, and please see Appendix B for a copy of a sample PODS input file) First, the system level inputs that control the settings for parameters affecting the simulation itself are presented. Next, the airline level inputs that determine the settings for the two competing airlines in the simulation are detailed. Finally, for the market inputs, the three networks used for all the experiments are described. Furthermore, the three testing sequences used will also be discussed. This should provide the reader with a good understanding of the basis for which all experiments were run under, which is essential before discussing the details of the results.

3.1 System Level Inputs

PODS system level inputs enable the control of the main settings of the simulation. For all simulations, the input parameters for the system level inputs were kept constant after some sensitivity analysis to determine their sensitivity and their appropriate values. This section details the system level inputs.

3.1.1 Top Level Settings

To simplify the experiments, a few top-level assumptions were made. In the PODS experimental world, there were only two airlines, competing with each other. There were three different network scenarios the two airlines compete in, and they will be discussed

in Section 3.3. To speed up the simulation as well as simplify the experiments, there were no no-shows and cancellations, and the airlines also did not overbook.

All passengers in the PODS world are grouped into two categories - business travelers and leisure travelers; and their characteristics are totally different. Leisure travelers are more sensitive to price and more flexible with their travel time, and since lower fares had advance purchase requirements, leisure travelers generally book early in order to get the cheaper fare. On the other hand, business travelers tend to book closer to the travel date because they are less sensitive to price, and may not know that they have to travel until just days before the travel date. Please see Figure 3.1 for the booking curves used in the experiments. As in the real world, in PODS most bookings for leisure travelers occur earlier, while most business travelers book within a week of the travel date.

Booking Curves for Business and Leisure Passengers




Figure 3.1: Booking Curves for Business and Leisure Passengers

The two types of passengers also have a different willingness-to-pay level, as well as different disutility costs for different restrictions. The willingness-to-pay parameter will

be discussed in experiments related to the PODS7 version of PODS in the next chapter, since we introduced a new representation for maximum willingness-to-pay in PODS7. As for disutility costs, it will be discussed in Section 3.2.1, which covers fares and fares structures, as well as the details for different disutilities in the PODS world.

3.1.2 Amount of Data Collected

The number of data points taken for each run was determined by three parameters: *ntrial* (number of trials), *nsamp* (total number of samples in each trial), and *nburn* (number of samples burned or thrown out in each trial before statistics were collected). Each PODS run contains several trials, and for each trial several hundred samples are taken, with a number of samples thrown out at the beginning to overcome the initial condition effects. Due to the nature of passenger demand and airline scheduling, the only comparable data is the same flight for the same day in different weeks. Therefore, each sample point taken in PODS represents a certain day of the week, so 52 data points from PODS might represent one year's worth of data from a network of flights every Tuesday afternoon, for example.

In order to find the best combination for statistical stability and minimum run time, an excursion was performed to determine the number of trials and number of observations to run. Network 1 (please see Section 3.3.1 for a detailed description of Network 1) was used, and the two airlines in this experiment both used EMSRb as their YM method. Since the two airlines were identical, we should be able to see a statistically insignificant difference in revenue in the simulation results. A confidence interval should be able to be constructed because the number of samples as well as the distribution of the data was known at the end of each trial. However, we have to note that the generic assumption of independence did not hold here. In order to construct any statistical interval, we had to make the assumption that all the data points were independent of each other. However, in PODS' case, due to the fact that each forecast incorporated observations from the previous 26 data points, the data points were not independent of each other. Therefore this experiment was necessary in order to determine the appropriate combination of the

number of trials to run and the number of samples in each trial. The results are presented here:

Case	# trials	# sample per trial	# sample burned	total # of samples	YM Methods		Revenue		Revenue Difference	% Diff in Revenue
					Airline A	Airline B	Airline A	Airline B		
Base	10	600	200	4000	EMSRb	EMSRb	199441	199911	-470	-0.24%
1	5	1000	200	4000	EMSRb	EMSRb	199525	199956	-431	-0.22%
2	20	1000	200	16000	EMSRb	EMSRb	199964	199957	7	0.00%
3	10	1000	200	8000	EMSRb	EMSRb	199873	199844	29	0.01%
4	20	600	200	8000	EMSRb	EMSRb	199803	199814	-11	-0.01%

Table 3.1: Results from the number of trials, number of samples per trial, and number of samples burned experiment

After experimenting with different trial and sample sizes, the conclusion was to use 20 trials, 600 samples, and 200 burns for each experiment. This led to $20 * (600-200) = 8000$ days, or 8000 samples for each experiment. From the test case above, the difference in revenue was only 0.01%, so this was a satisfactory combination of trials, sample sizes, and burns.

3.1.3 Demand Variability Parameters

Finally, the ability to control the variability of demand is one of the features that sets PODS apart from more traditional simulations. There were three parameters that control the demand variability: The system k-factor, the market k-factor, the passenger type k-factor. K factor is defined as:

$$K \text{ factor} = \sigma/\mu$$

where σ = standard deviation

$$\mu = \text{mean}$$

The system K factor represents passenger demand variation for the whole system. A system variation multiplier is first generated from the K factor value, which later is combined with other multipliers to modify the passenger demand. The value selected for

system K factor was 0.1, meaning that the multiplier value was randomly selected with a mean of 1 and a standard deviation of 0.1.

The market K factor represents passenger demand variation among different markets. The value used was 0.2, so a random variate multiplier was generated with a mean of 1 and a standard deviation of 0.2.

The passenger type K factor represents passenger type demand variation. The demand for the two passenger types in PODS simulations are input separately, so each passenger type gets their own passenger type variation multiplier. The value used for passenger type K factor was 0.4, so for each type of passenger in each market, a random variate multiplier with a standard deviation of 0.4 was generated.

These first three K factors – the system, the market, and the passenger type K factors, generates three different multipliers for each passenger type in each market. The input demand is then multiplied by the combined multiplier to generate the final demand used in each simulation run.

3.1.4 Disutility Cost Parameters

Similar to the variation in demand, the variation in disutility cost is controlled by the attributed cost K factor. A setting of 0.3 was used, meaning that with a disutility value set to 0.9 for a particular restriction, a traveler would perceive it as a disutility of $(basefare) * 0.9$ dollars, with a standard deviation of $(basefare) * 0.9 * 0.3$ dollars.

The disutilities for each type of passenger are then combined with the fares in each O-D market to give passengers a perceived total cost for each fare choice. The passengers then pick the fare product with the lowest total combined cost. For example, in a 1000 mile market, the *basefare* was \$200. Q class carried 3 restrictions with it, and for a leisure passenger, the attributed cost in terms of *basefare* was $1.75 + 0.25 + 0.25 = 2.25$. As a result, the expected perceived value for Q class to leisure travelers was then $\$200 +$

$2.25 * basefare$. Since $basefare = \$200$, the perceived cost was then \$650. After adding in other costs in a similar fashion, such as replanning and cost assigned to the unfavorable airline, this value is then compared with other available fare products as well as that particular passenger's maximum willingness-to-pay value. Finally that passenger then chooses the path with the least perceived cost that also fell below their maximum willingness-to-pay value.

In summary, the system level inputs determine the main settings for the PODS simulation. For these inputs, assumptions were made or sensitivity tests were done in order to come up with a suitable value, then for all the PODS experiments the same values were used for consistency.

3.2 Airline Level Inputs

This section describes the input parameters that determine the behavior or decision of each airline. Presented here are the fare classes each airline offered and their restrictions, as well as the parameters for the revenue management system. For each of the two airlines in the PODS simulation, a different fare structure as well as different revenue management selection can be made in order to explore the competitive effects.

3.2.1 Fare Class Structures

To simplify the experiments, both airlines offered 4 fare classes in all O-D markets: Y, B, M, and Q class. The differences between the four classes were their fares, their advance purchase requirements, and their restrictions. The fare structure in the PODS world was set so that it approximated the fare structure of many US airline markets today. The Q class fare was the "*basefare*", and Y class fare was 4 times the *basefare* value, followed by B class at double the *basefare*, and finally M class at 1.5 times the *basefare*. For advanced purchase requirements, Y class had no restrictions just like the unrestricted full fare on the market today, followed by M class with a 7-day advance purchase requirement, B class with a 14-day advance purchase requirement, and finally Q class

with a 21-day advance purchase requirement. As for the restrictions, 3 different restrictions were designated according to industry practice: Saturday night stay, cancel/change penalty, and non-refundability. B class had the Saturday night stay restriction, M class added on the cancel/change penalty, while Q class had those two restrictions plus non-refundability.

For the two types of passengers, each type had a different disutility cost assigned to each restriction. For business passengers, since they generally do not want to be on the road on the weekend, their disutility cost for Saturday night stay was $2.25 * \text{basefare}$. As for cancel/change penalty and non-refundability, their disutility costs were both $0.75 * \text{basefare}$. After adding in the ticket cost, this caused an increasing level of disutility, thus making the average business traveler prefer the higher fare Y class, all else being equal. For leisure travelers, the opposite was true. They are more sensitive to price, and are more willing to spend over a Saturday night in order to get a lower fare. Therefore, their disutility cost for Saturday night stay was $1.75 * \text{basefare}$, lower than business traveler's value of $2.25 * \text{basefare}$. Their disutilities for cancel/change penalty and non-refundability were $0.25 * \text{basefare}$, both also lower than business traveler's values. After adding in fares, the resulting decreasing disutility made the average leisure traveler prefer the lowest, restricted Q class, all else being equal.

For an example market of 1000 miles in length, the following table shows its fares, disutility costs, and finally perceived combined cost values for both business travelers and leisure travelers.

Fare Class	Fare	Advance Purchase	Restrictions	Attributed Cost in Terms of <i>basefare</i>		Perceived Cost	
				Business	Leisure	Business	Leisure
Y	\$800	0	None	0	0	\$800	\$800
B	\$400	7 days	1	2.25	1.75	\$850	\$750
M	\$300	14 days	1, 2	$2.25 + 0.75$	$1.75 + 0.25$	\$900	\$700
Q (<i>basefare</i>)	\$200	21 days	1, 2, 3	$2.25 + 0.75 + 0.75$	$1.75 + 0.25 + 0.25$	\$950	\$650
Note:	Restriction 1: Saturday Night Stay						
	Restriction 2: Change/Cancel Penalty						
	Restriction 3: Non-refundability						

Table 3.2: Fares, disutility costs, and perceived fares for business and leisure travelers in a 1000 mile market

As can be seen from the example above, the way the restrictions and attributed costs are set-up in PODS, the average business traveler perceives an increase in total cost from Y to Q class, even though the actual fares decrease. This made them favor the Y class over B, M, and Q class. Similarly, leisure travelers saw a decrease in perceived fares from Y to Q class, thus making them favor Q class. One note of caution, however, since the attributed costs have a K factor of 0.3, the perceived fares shown above are expected values. Actual values would be different for each passenger, since the attributed costs vary randomly around the mean value reported in the example.

With the above fare structures as well as disutility costs in place, the next section in the PODS simulation input is the revenue management system.

3.2.2 Revenue Management System

This section is where the most numerous changes in input parameters were made through the experiments. The Revenue Management input section controls the RM optimization method used, virtual class definitions, forecasting and detrunctation methods, as well as availability processor controls. For RM optimization methods used, the choices were First-Come-First-Served (FCFS), EMSRa, EMSRb, GVN, DAVN, Network Bid Price (Netbid), and finally Heuristic Bid Price (HBP). Chapter 2 already gave a description of how each algorithm works. In this section, specific implementations of each algorithm, especially the options available and used in PODS experiment runs, will be discussed.

EMSRb: EMSRb fare class control was the base case, and it was a leg-based algorithm. Airlines uses demand-weighted total Origin-Destination Fares (ODF) as revenue inputs to calculate the EMSRb booking limits.

GVN: GVN is still a leg-based algorithm, but the fares for the whole network were divided into six virtual buckets. Please see the network descriptions below for the specific virtual class boundaries for each network used. ODFs are assigned to virtual buckets based on their total fare value, and EMSRb algorithm is used to calculate the virtual class booking limits.

Netbid: Netbid is short for deterministic network bid-price optimization model, which uses ODF demand forecasts to determine leg shadow prices for all legs traversed by each ODF. In the PODS experiments using Netbid, a booking is accepted if the fare is higher than the sum of all shadow prices for all the legs the passenger wanted to travel. Since there is in essence no “booking limit” with Netbid, this could lead to the airline accepting too many low yield passengers between bid-price re-optimizations. As long as these passengers pay a price higher than the sum of all shadow prices, all of them would be accepted. This led initially to poor performance for Netbid, so two new parameters were introduced in PODS to improve the algorithm.

The first of the new parameters was *nreop*, which determined the number of bookings between re-optimization. Since one drawback for Netbid was that it accepts any bookings that had a higher fare than the sum of the shadow prices, we needed to come up with a “cap” so that the airline using Netbid would not be flooded with low fare passengers. *nreop* determines how often the network algorithm was executed to determine the booking limits. No bookings would be taken if a particular fare class was closed. The setting used for all PODS experiments for *nreop* was 10, meaning that after every 10 bookings, the re-optimization process was repeated.

The second new parameter was *iblv*, which stands for “I-Believe”. This parameter determines whether the forecast made by PODS are to be “believed” by the revenue

management system or not. If the forecast made was “believed”, then the forecast was decremented by 1 each time a booking was made. The setting used for all PODS experiments was that the forecast was “believed”, and it was decremented by 1 when a booking was made.

DAVN: DAVN stands for Displacement Adjusted Virtual Nesting. The assignment of ODFs to virtual buckets is based on “pseudo-fares” instead of ODF total fare values. The pseudo-fares for a 2-leg itinerary are calculated as follows:

Pseudo-fare for leg 1 = total fare – shadow price for leg 2

Pseudo-fare for leg 2 = total fare – shadow price for leg 1

Finally for each virtual class, its booking limits are calculated using EMSRb.

PODS added two capabilities for DAVN. First was the ability to set the virtual class boundaries so that the demands were equalized among virtual classes. For most experiments in this thesis, the virtual class boundaries were the values described in the Network descriptions in Section 3.3. In Section 4.5.3, experiments were made to introduce the demand equalized virtual class boundaries, and the results will be discussed there. With the demand equalized virtual class boundaries, the virtual class boundaries were set by making the demands in every virtual class roughly equal. The result of this change can be seen in Section 4.5.3.

The second capability specific to DAVN in PODS was the ability to control whether the network displacement costs and the boundaries were calculated once at the beginning of each departure’s booking process, or at every time frame. The default value used for most of the experiments here was only to calculate them once at the beginning of each departure’s booking process, or each sample. As described in Section 4.5.3, another experiment was done to investigate the effect of re-calculating the network displacement costs and the boundaries at every time frame. Combined with the different virtual class boundary schemes described above, several different implementations of DAVN were tested.

HBP: HBP is Heuristic Bid-Price. For local itineraries, HBP treats local legs exactly the same as GVN. For connecting itineraries, a 2 leg itinerary is accepted if its fare is greater than $\max[\text{EMSR1}, \text{EMSR2}] + \text{d-factor} * \min[\text{EMSR1}, \text{EMSR2}]$ where EMSR1 is the EMSRb value at the capacity of leg 1, and EMSR2 is the EMSRb value of leg 2. D-factor was selected by the operator for different local-connection passenger mix, and a value of 0.25 was used for all PODS experiments in this thesis. Different D-factors were tested as well, and the results of that experiment is presented in Section 4.5.2.

Similar to Netbid, HBP accepts a connection booking if its fare value is greater than the value calculated using the formula given above. For connecting traffic, there is no booking limit to prevent too many seats taken up by low yield passengers. In order to overcome this problem, again *iblv* and *nreop* were introduced. Again *iblv* was set so that the forecast was “believed”, and the forecast was decremented by one with a booking in that fare class. For *nreop* however, we decided to use a setting that no re-optimizations occur. HBP seems to be more robust, and the formula above seems to prevent too many low fare bookings, so frequent re-optimization is not needed.

As for forecasting and detruncation for the Revenue Management system, PODS has many different methods available to choose from. For forecasting, the methods includes pick-up forecasting and regression forecasting. For detruncation, the methods include probabilistic booking curve detruncation and projection detruncation. For a detailed description of the functionality for each method, please see Zickus (1998).

The airline level inputs included parameters to control the fare structures as well as the Revenue Management systems. The majority of the experiments in this thesis involved changing characteristics of the revenue management systems or passenger behavior, both of which were covered in this section.

3.3 Market / Network Level Inputs

Three different networks were created for the PODS experiments presented in this thesis. For all three networks, the two airlines each operate a hub operation to various spoke cities. The flight times and capacity offered by each airline in each market were identical. A PODS day consisted of one connecting bank, where for each airline, flights from spoke cities fly into the hub city then return to the spoke city. Both airlines offer four fare classes in all their O-D markets – Y, B, M, and Q classes, and they matched each other's fares in all markets.

3.3.1 Network 1

Network 1 consisted of one hub city and four spoke cities. Both airlines operated out of the same hub city, thus competing with each other for both local and connecting traffic. This network was similar to the situation at Chicago O'Hare, where two major airlines co-located their hubs at the same airport and competed for both local and connecting traffic.

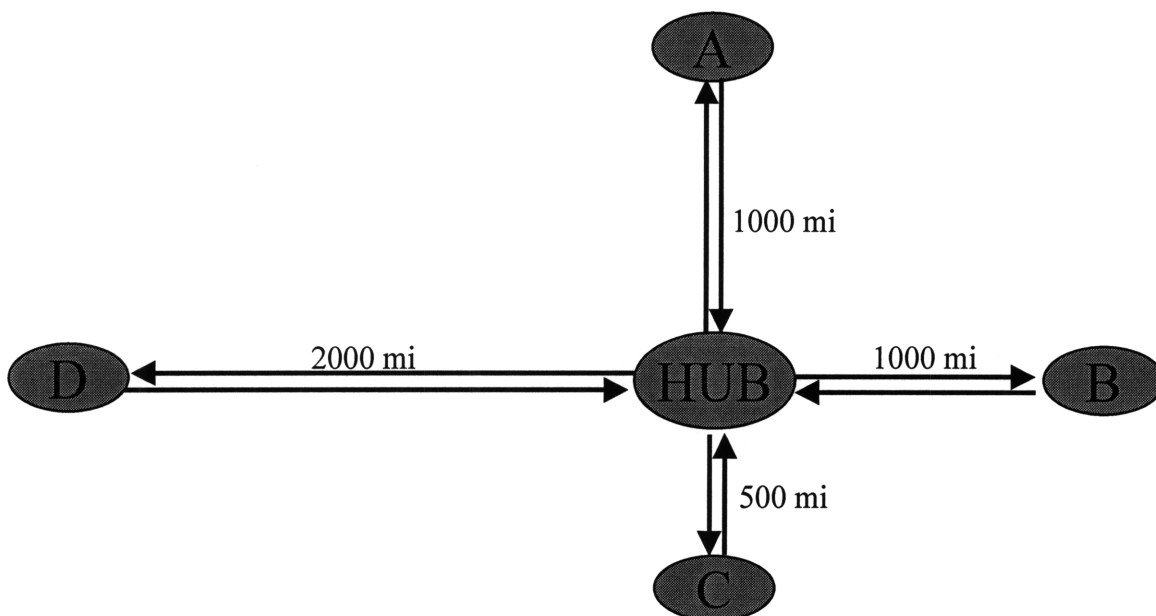


Figure 3.2: Network 1 - Two airlines co-locate their hubs, competing for both local and connecting traffic

Network 1 consisted of one short-haul leg, two medium-haul legs, and one long-haul leg. The demands were designed so that the short-haul leg would have the highest demand, while the long-haul leg would have the lowest demand. This was done to simulate the real world situation where feeder flights from close-in cities to the hub generally have higher demand than transcontinental flights. The following is a summary of leg distances and demands for network 1. Note that the demand shown in the table below was demand for each airline, in the market indicated. All aircraft had a capacity of 100 seats.

City	Distance from Hub
A	1000
B	1000
C	500
D	2000

Table 3.3: Network 1 Leg Distances

O-D Market	Mean Demand	Leg	Mean Leg Load	Local %
A-H	48	A-H	112	43%
H-A	48	H-A	112	43%
B-H	32	B-H	96	33%
H-B	32	H-B	96	33%
C-H	60	C-H	132	45%
H-C	60	H-C	132	45%
D-H	16	D-H	72	22%
H-D	16	H-D	72	22%
A-B	24	Average	103	36%
A-C	24			
A-D	16			
B-A	24			
B-C	24			
B-D	16			
C-A	24			
C-B	24			
C-D	24			
D-A	16			
D-B	16			
D-C	24			

Table 3.4: Demand Summary for Network 1

Note that the demand shown here is for a demand factor of 1.0. The goal was to have the average demand on all legs equal the aircraft capacity of 100. Since the resulting average of 103 was very close to 100, this demand scenario is still considered DF = 1.0. For other

demand factors, such as $DF = 1.2$, all the demands listed here are increased by 20%. Likewise for $DF = 0.8$, demands are decreased by 20%.

3.3.2 Network 2

Network 2 consisted of two hub cities and again four spoke cities. Each airline operated out of their own hub cities. They also offered services to the same four spoke cities, thus competing with each other only for connecting traffic but not for local traffic. This network was similar to other hub cities in the domestic United States, where major airlines have their own hubs but they also serve the same spoke cities.

The leg distances as well as demand values for network 2 were the same as for network 1. The two hub cities both happen to be at the same distance from all the spoke cities. Again shorter haul legs had a higher demand than the longer haul legs.




Figure 3.3: Network 2 – Two airlines, each with their own hubs, competing only for O-D traffic

3.3.3 Network 3

Network 3 consisted of two hub cities and six spoke cities. During the investigation for Revenue Management method Network Bid Price (Netbid), it was suspected that some of our results were due to the small size of the simulation network. One solution to this problem was to increase the number of spoke cities from four to six, thus almost doubling the number of markets from 28 to 54. Again, each airline operated out of their own hub and they also offered services to the same six spoke cities, thus competing with each other only for connecting traffic but not for local traffic.

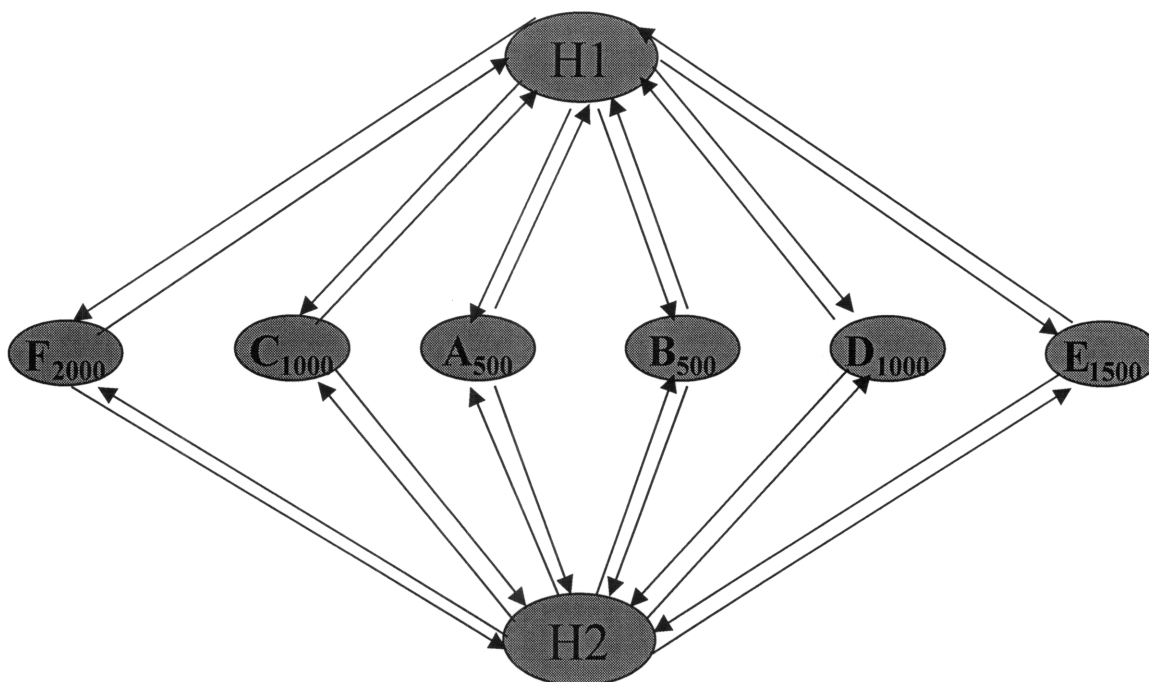


Figure 3.4: Network 3 - Two airlines with separate hubs, competing only for connecting traffic

Similar to Network 2 except with 6 cities instead of 4, Network 3 had two short-haul, two medium-haul, and two long-haul legs. Again the demand for the short haul legs was higher than that for longer haul legs.

City	Distance from Hub
A	500
B	500
C	1000
D	1000
E	1500
F	2000

Table 3.5: Network 3 Leg Distances

O-D Market	Mean Demand	Leg	Mean Leg Load	Local %
A-H	65	A-H	140	46%
H-A	65	H-A	140	46%
B-H	50	B-H	125	40%
H-B	50	H-B	125	40%
C-H	45	C-H	110	41%
H-C	45	H-C	110	41%
D-H	35	D-H	90	39%
H-D	35	H-D	90	39%
E-H	25	E-H	75	33%
H-E	25	H-E	75	33%
F-H	20	F-H	60	33%
H-F	20	H-F	60	33%
A-B	21	Average	100	40%
A-C	17			
A-D	14			
A-E	13			
A-F	10			
B-A	19			
B-C	18			
B-D	15			
B-E	13			
B-F	10			
C-A	17			
C-B	16			
C-D	13			
C-E	11			
C-F	8			
D-A	15			
D-B	14			
D-C	11			
D-E	9			
D-F	6			
E-A	13			
E-B	13			
E-C	11			
E-D	7			
E-F	6			
F-A	11			
F-B	11			
F-C	8			
F-D	6			
F-E	4			

Table 3.6: Demand Summary for Network 3

3.3.4 Parameters Common to the Three Networks

The fares and the virtual class boundaries were identical in the three networks, and they are presented here. The fares for the Y class were 4 times the *basefare*; the fares for the B class were 2 times the *basefare*; the fares for M class were 1.5 times the *basefare*; finally the fares for Q class were set to be equal to the *basefare* value input to PODS. The fare levels for each O-D market were set based on a distance formula, where doubling the distance travel equaled roughly 1.6 times the fare. All the fare presented here are one-way fares.

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

Table 3.7: Fare Chart for PODS Networks

As for virtual class boundaries, 6 virtual classes were used for the simulations of VEMSRb, DAVN, and HBP (other than those discussed in Section 4.5). They were manually selected so that each class had a good mix of Y, B, M, and Q fares, and the four fares for each O-D market were mostly distributed in different virtual classes to provide separation.

Virtual Class #	Fares (\$)	Virtual Class Upper Boundary
1	1280	1300
	1205	
	1020	
	920	
	800	
2	660	700
	640	
	602.5	
	510	
	500	
	480	
	460	
451.88		
		420

3	400 382.5 345 330 320	
4	301.25 300 255 250 247.5 230	310
5	200 187.5 165 160	215
6	125 120 80	140

Table 3.8: Virtual Classes Ranges

In this section, the three networks used for the PODS simulations were presented. Network 1 was a simple, co-located hub with 4 spoke cities. The two airlines competed with each other for both local and connecting traffic. Network 2 was a variation of network 1, where the two hubs were not co-located, and the airlines competed only for connecting traffic. Finally Network 3 added two cities onto Network 2 thus doubling the number of markets.

3.4 PODS Testing Sequences

There were three sequences used in PODS experiments – Sequence 0, Sequence 2, and Sequence 3. This Section describes the three sequences used for all experiments.

3.4.1 Sequence 0

In order to validate the results from PODS and compare them to other traditional simulations, such as the MITSIM, Sequence 0 was created. Since MITSIM had the assumption of independent ODF demands and demand arrival being a Poisson process,

PODS settings in Sequence 0 were set up to replicate them. For the assumption of independent ODFs, a parameter, *FCOC* (First-Choice-Only-Choice), was introduced. In Sequence 0, this parameter was set so that it implied perfect booking history as if the RM system had a perfect crystal ball that could see future bookings. It also made a passenger's choice their only choice. This sequence 0 setting ensured that the path/fare demand was fixed, and was independent of availability.

Furthermore, in Sequence 0 the system, market, and passenger type K-factors were set to zero. The Demand Z-factors were set to 1.0. All these combined together ensured that the assumption in MITSIM of ODFs being independent was repeated. This allowed a direct comparison between the MITSIM and PODS, and allowed us to check the validity of PODS results.

3.4.2 Sequence 2

Sequence 2 was similar to Sequence 0 because it still assumed First-Choice-Only-Choice for all passengers. However, in order to investigate what the effect demand variation had on revenue outcome, the standard system, market, and type K-factors were used. As a result, Sequence 2 still assumed that path/fare class demand was fixed and independent of availability, but ODF demands were correlated and more variable than in Sequence 0.¹

3.4.3 Sequence 3

Sequence 3 was the most used Sequence in PODS experiments; it was also the one that resembled the real world the best and explored all PODS' capabilities to its fullest. In Sequence 3, every passenger had more than one choice as to which itinerary they want to travel. This meant that passengers could sell-up to a higher fare class or fly on the competitor if their first-choice fare class was unavailable. As for demand variation, all the PODS standard K and Z factors were used, as described in Section 3.1.1. Sequence 3 was used most often because it simulated higher demand variation as well as passenger choice. The ability to do so was what set PODS apart from traditional simulations.

¹ Taken from PODS Summit IV Report, MIT Flight Transportation Laboratory, July 1997

The three sequences described in this Section covered a wide range of possible combination of independence and variation of demand and fare class. We started in Sequence 0 with independent demand and no demand variance, followed by Sequence 2 with demand variance, finally Sequence 3 with full demand variance and correlated ODF demands and dependent fare class availability. The three Sequences allowed us to investigate the effects of having different travel choices, as well as the variation in demand.

3.5 Chapter Summary

The PODS simulation environment was discussed in this chapter. All the input parameters, including system level, airline level, and market level inputs, were presented. The airline revenue management input parameters, as well as the three networks used for all the PODS studies were described in detail. With the “PODS world” described, the reader should be ready to move into the PODS experiments and results.

Chapter 4

Results and Discussion

In this chapter, the questions raised in Chapter One will be discussed, and the simulation results will be used to answer these questions. The questions covered in this chapter include:

- 1) The relative performance of network Origin-Destination control Revenue Management methodology compared to EMSRb fare class control methodology -- does network Origin-Destination control increase revenues, and which methods perform best?
- 2) Impacts of passenger choice assumptions -- how do revenues and loads change under different passenger choice assumptions?
- 3) Competing airlines with co-located hub vs. separate hubs – what are the differences in loads and revenues?
- 4) Different implementations of DAVN and HBP O-D seat inventory control methods -- for DAVN and HBP, what are the effects of different implementations on passenger load and revenue?

In this Chapter, the above four questions will be answered and discussed as a way to present the PODS O-D research results.

4.1 The Relative Performance of Network Origin-Destination Control Revenue Management Methodology Compared to EMSRb Fare Class Control Methodology

4.1.1 Motivation

As airlines around the world start implementing Network Origin-Destination Control for their revenue management system, it is important to first investigate whether the network O-D control Revenue Management methods have an advantage over leg based EMSRb fare class control. Theoretically, we all know that local optimization does not lead to global optimization; thus over a network, leg based fare class control does not always lead to a network optimal solution. However, is this true under realistic data collection, forecasting, and passenger choice? Can we claim that Network O-D control Revenue Management methods perform better than leg based EMSRb method? The PODS simulation was used to check the revenue results with one airline using the now standard leg based EMSRb fare class control, with the other airline implementing other revenue management methods. The results in this section will tell us whether leg based fare class control or O-D control RM methods can produce higher revenue, and that whether it is a zero-sum game between competing airlines.

4.1.2 Simulation Setup

For this experiment, Network 1, Network 2, and Network 3 were used under PODS6. Again to refresh our memory, Network 1 consists of one common hub for both airlines, plus four spoke cities which both airlines serve. Network 1 is somewhat “virtual friendly”, meaning that short legs have the highest demand, while long legs have a lower demand. This leads to a better performance for virtual-bucket based RM methods because they give preference to long-haul connecting passengers who bring in higher revenue. Network 2 is very similar to Network 1, except that the two airlines each have their own hub, and they only compete for connecting traffic. Network 3 is a bigger version of Network 2, with 6 spoke cities instead of 4 spoke cities.

As for the Revenue Management methods compared in this experiment, for all cases, Airline B used EMSRb as a base case. Airline A, however, was simulated under different RM methods throughout the experiment so that we could see the effect of the base case vs. O-D control Revenue Management methods. The RM methods Airline A used included EMSRb, GVN, Netbid, DAVN, and HBP. (see Section 2.3 for details of each algorithm)

As for results reported here, all the percentage-change numbers refer to changes versus the base case, which was EMSRb vs. EMSRb. For example, if for the case of DAVN vs. EMSRb, we report a change in revenue of 1% for the airline using DAVN, and -1% for the airline using EMSRb, then that means the airline using DAVN received 1% more revenue than the base case, and the competing airline using EMSRb received in 1% less revenue than the base case.

4.1.3 Results and Discussions

In this section, the revenue results for average network Demand Factor = 0.8 will be presented first, followed by DF = 1.0 and DF = 1.2. Also, passenger loads on an example leg and in an example market will be presented in order to explain where the revenue advantages of the RM algorithms came from.

			Network 1		Network 2		Network 3	
			% increase compared to		% increase compared to		% increase compared to	
YM methods			EMSRb	EMSRb	EMSRb	EMSRb	EMSRb	EMSRb
DF	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B
0.8	EMSRb	EMSRb						
0.8	GVN	EMSRb	0.60%	-0.81%	0.33%	-0.30%	0.77%	-0.70%
0.8	Netbid	EMSRb	-0.77%	-0.73%	-0.05%	0.22%	0.01%	-0.54%
0.8	DAVN	EMSRb	1.14%	-0.55%	0.48%	-0.55%	1.14%	-0.66%
0.8	HBP	EMSRb	0.12%	-1.03%	0.75%	-0.30%	0.73%	-0.83%

Table 4.1: Revenue Summary, DF = 0.8, Networks 1, 2, and 3

The results for Demand Factor = 0.8 under the three networks are presented in Table 4.1. Again, DF = 0.8 means that the average demand on all the legs over the whole network is equal to 80% of the aircraft capacity of 100 seats. The revenue differences between

different networks will be discussed in Section 4.4. In this section, the focus will be on the relative performance of O-D based RM methods versus the base case, which was leg-based EMSRb vs. EMSRb.

As we can see from Table 4.1, DAVN had the best revenue performance under Network 1. The next best algorithm was GVN, followed by HBP, and finally Netbid. Under Network 2, HBP had the best revenue performance, followed by DAVN, GVN, and Netbid. Finally under Network 3, DAVN did the best, followed by GVN and HBP. Again Netbid performed the worst out of the four O-D RM algorithms. DAVN, GVN, and HBP consistently produced positive revenue results when compared to the base case which used the leg based EMSRb fare class control.

Netbid consistently performed worse than our base case under both Networks 1 and 2, which have 4 spoke cities each. The suspicion was that under a small network, the number of ODFs is smaller, thus the gaps between bid-prices were too large for Netbid. The airline using Netbid was not able to control low-fare bookings effectively because it accepted bookings by comparing fares to bid-prices, leading to its poor revenue performance. In the larger Network 3 with 6 spoke cities, there are many more ODFs compared to Networks 1 and 2. This increased number of ODFs resulted in smaller steps between bid-prices, and Netbid could then control low fare bookings more effectively. Netbid performed at least as well as our base case leg based EMSRb under Network 3, supporting our suspicion that the small-network effect lead to Netbid's poor performance. However, even with the 6-city network, Netbid stilled performed the worst out of the four O-D control RM methods. As it turned out from investigations by Zickus¹, another reason for the poor performance by the airlines using Netbid was that Netbid simply performed poorly with the basic pick-up forecasting and booking curve detruncation used in our experiments. Netbid actually performed as well as the other O-D control methods under different forecasting and detruncation methods. For a detailed experimental result

¹ Zickus, J.S. *Forecasting for Airline Network Revenue Management: Revenue and Competitive Impacts*, June 1998, Section 4.1.

for Netbid under different forecasting and detrunctation methods, please see Zickus (1998).

Another interesting point with the $DF = 0.8$ results above was that we are seeing some “zero-sumness”. When the airline using an O-D control RM method had a revenue gain compared to the base case, the competing airline using EMSRb fare class control in most cases had a revenue loss. This indicated that part of the revenue gain of the first airline came at the expense of the second airline. This also indicated that by implementing an O-D control method, airlines can take revenue away from competitors that are still using leg based fare class control. The revenue gain for an airline not only comes from getting more money from its own passengers, but also from taking away some of the competitor’s high-revenue passengers.

Moving onto Demand Factor of 1.0, where Table 4.2 shows the revenue results for the three networks. At this higher level of demand, the results are still very similar to the results under $DF = 0.8$.

			Network 1		Network 2		Network 3	
			% increase		% increase		% increase	
			compared to		compared to		compared to	
	YM methods		EMSRb	EMSRb	EMSRb	EMSRb	EMSRb	EMSRb
DF	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B
1.0	EMSRb	EMSRb						
1.0	GVN	EMSRb	0.37%	-0.30%	0.26%	-1.58%	0.56%	-1.91%
1.0	Netbid	EMSRb	-0.64%	0.08%	-0.66%	-1.05%	0.52%	-1.38%
1.0	DAVN	EMSRb	1.92%	-1.07%	1.46%	-2.03%	2.09%	-1.92%
1.0	HBP	EMSRb	1.71%	-0.68%	1.32%	-1.78%	1.77%	-1.69%

Table 4.2: Revenue Summary, $DF = 1.0$, Networks 1, 2, and 3

For both Networks 1 and 2, DAVN performed the best, followed by HBP and then GVN. Netbid again performed poorly in the small network by showing a negative revenue impact. For Network 3, the relative performances of the Revenue Management methods stayed the same. One interesting thing to note is that Netbid under the bigger network actually produced a positive revenue gain over the base case, confirming the theory that part of the reason for its bad performance under the smaller networks were due to the

small network effect. The airline using Netbid still performed worse than the airlines using other O-D RM methods due in part to the pick-up forecasting and booking curve detruncation methods used.

As for zero-sumness, again at this higher level of demand the gains in revenue by the airline using O-D control methods was accompanied by a loss in revenue by the airline using EMSRb fare class control. This again suggests that by implementing an O-D control method, the gain in revenue came partly at the expense of the competitor.

Finally at an even higher level of demand at Demand Factor of 1.2, a similar yet somewhat different result emerged. Table 4.3 shows the result under all three Networks at DF = 1.2.

			Network 1		Network 2		Network 3	
			% increase		% increase		% increase	
			compared to		compared to		compared to	
YM methods			EMSRb	EMSRb	EMSRb	EMSRb	EMSRb	EMSRb
DF	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B
1.2	EMSRb	EMSRb						
1.2	GVN	EMSRb	-0.84%	0.65%	0.46%	-0.99%	1.23%	-1.33%
1.2	Netbid	EMSRb	-0.94%	0.63%	-0.47%	-0.58%	1.08%	-1.06%
1.2	DAVN	EMSRb	2.68%	-0.45%	2.44%	-1.67%	2.60%	-1.63%
1.2	HBP	EMSRb	1.91%	-0.14%	2.47%	-1.56%	2.41%	-1.51%

Table 4.3: Revenue Summary, DF = 1.2, Networks 1, 2, and 3

Under Networks 1 and 3, DAVN performed the best, followed by HBP, GVN, and finally Netbid. Under Network 2, DAVN and HBP performed about the same, with HBP performing slightly ahead of DAVN. This is contrary to the previous findings at lower demand factors where DAVN performed the best under all conditions. However in general, under all three Networks for all three demand factors, DAVN still performed the best, followed closely by HBP. GVN was the next best performer, and finally Netbid again resulted in a revenue loss compared to the base case under Networks 1 and 2. In Network 3, even though Netbid was able to produce a positive revenue gain, its performance still lagged behind other O-D control RM methods.

GVN under Network 1 at $DF = 1.2$ showed a negative result (-0.84%), while it still showed positive results under Networks 2 and 3. This was because under Network 1 at a high demand factor, GVN became too greedy and took on too many connecting passengers. These connecting passengers end up displacing high yield local passengers, who were then spilled to the competitor. This effect can be seen by the 0.65% revenue gain by the competitor using EMSRb under Network 1. Furthermore, under Networks 2 and 3, since the airlines do not compete for local passengers, the greediness of GVN does not result in spilling the local passengers to the competitor, and thus GVN results remained positive. For a further discussion on the difference between Networks 1 and 2 as well as the effect of de-coupling the hubs, please see Section 4.4.

As for zero-sumness, again under $DF = 1.2$ we see strong evidence that the revenue gain of the airline practicing O-D control came at the expense of the airline practicing leg fare class control. This again confirmed that at all levels of demand, using O-D control RM methods results in a competitive advantage over leg fare class control.

After establishing that O-D control RM methods perform better than EMSRb fare class control, in this next section we will try to find out where the revenue advantage came from by looking at the leg class distribution as well as path class distribution. Network 2 was used because it is a smaller network than Network 3, making analysis easier. Furthermore, its de-coupled hubs are more common in the real world than the co-located hub situation in Network 1. A demand factor of 1.0 was used for this analysis.

First, we consider the EMSRb vs. EMSRb leg load results. EMSRb fare class control was used as our base case for the comparisons. Note that three sets of leg demand results are presented – High, Medium, and Low. These three results were passenger loads on three example legs in the network, not the aggregated results of several legs.

DF	Airline A	Airline B	Leg	Airline A					Leg	Airline B				
				Passenger Load by Class						Passenger Load by Class				
				Total	Y	B	M	Q		Total	Y	B	M	Q
1.0	EMSRb	EMSRb	High	90.91	43.14	14.88	20.10	12.79	High	90.96	43.05	14.84	20.22	12.85
					47%	16%	22%	14%			47%	16%	22%	14%
			Medium	88.00	36.75	12.84	10.32	28.09	Medium	87.90	36.47	12.92	10.14	28.37
					42%	15%	12%	32%			41%	15%	12%	32%
			Low	60.11	24.20	8.80	6.30	20.81	Low	60.08	24.35	8.82	6.21	20.70
					40%	15%	10%	35%			41%	15%	10%	34%

Table 4.4: EMSRb vs. EMSRb Leg Passenger Load by Class Results

In Table 4.4, the base case EMSRb vs. EMSRb passenger load by class results are presented. On an aircraft with 100 seats, the high demand leg resulted in an average load of more than 90 passengers, while the low demand leg resulted in only about 60 passengers. This result will be used for comparison to investigate how the O-D control results differed from the base case results.

Next, we examine the GVN vs. EMSRb leg load results. GVN's performance was not as good as DAVN or HBP, but was better than Netbid. In Table 4.5 the passenger load by class results are presented.

DF	Airline A	Airline B	Leg	Airline A					Leg	Airline B				
				Passenger Load by Class						Passenger Load by Class				
				Total	Y	B	M	Q		Total	Y	B	M	Q
1.0	GVN	EMSRb	High	90.94	42.98	12.68	7.77	27.51	High	89.58	43.33	14.98	13.72	17.55
	compared to base				47%	14%	9%	30%			48%	17%	15%	20%
	0.26%	-1.58%	Medium	87.94	36.13	12.57	11.25	27.99	Medium	86.76	36.68	12.93	8.12	29.03
					41%	14%	13%	32%			42%	15%	9%	33%
			Low	70.11	23.72	9.07	4.68	32.64	Low	56.28	24.07	8.58	3.10	20.53
					34%	13%	7%	47%			43%	15%	6%	36%

Table 4.5: GVN vs. EMSRb Leg Passenger Load by Class Results

For the particular case presented above, the airline using GVN had a revenue gain over the base case of 0.26%, while the competing airline using EMSRb had a revenue loss of 1.58%. The airline using GVN saw a huge increase in passenger load for the leg with low demand, from 60.11 to 70.11. Most of that increase in passenger load came in its Q class, which went from 20.81 to 32.64 for that particular leg. By carrying about 10 more passengers on the low demand leg while carrying about the same number of passengers on the high and medium legs, GVN was able to improve its revenue. On the other hand, on the high demand leg, even though the total load was about the same, the airline using

GVN actually had a worse fare class distribution. It carried 27.51 Q class passengers, versus 12.79 in the base case. It also carried only 7.77 M class passengers, which was much worse than the base case when 20.10 M passengers were carried. This lower-yield fare class distribution had less effect on revenue than the 10 extra Q passengers carried on the low leg, thus giving the GVN airline a 0.26% improvement in revenue.

As for the airline that still used EMSRb fare class control, they suffered a 1.58% loss in revenue. Compared to the base case where both airlines used EMSRb, this second airline carried fewer passengers on all legs. Its passenger fare class mix was also worse, carrying fewer passengers, especially in M class. These lost passengers were probably carried by the first airline using GVN because the GVN airline carried more passengers on most legs than the base case. This again suggests the zero-sumness effect which states part of the gain for the airline using O-D control RM methods is at the expense of the competing airline still using EMSRb fare class control, and this shift in passengers carried between the two airlines confirms the zero-sumness effect.

Moving on to Netbid vs. EMSRb leg load results, Netbid performed poorly under most cases, and its leg load results can explain why. Table 4.6 shows the passenger load by class results for Netbid vs. EMSRb.

DF	Airline A								Airline B					
	Airline A	Airline B	Leg	Total	Y	B	M	Q	Leg	Total	Y	B	M	Q
1.0	Netbid	EMSRb	High	92.98	41.55	12.32	6.93	32.18	High	89.65	43.75	15.37	16.12	14.41
	compared to base				45%	13%	7%	35%			49%	17%	18%	16%
	-0.66%	-1.05%	Medium	89.42	35.22	12.00	5.07	37.13	Medium	87.34	37.02	13.22	10.20	26.90
					39%	13%	6%	42%			42%	15%	12%	31%
			Low	70.74	23.62	9.02	4.54	33.56	Low	55.64	24.35	8.72	3.14	19.43
					33%	13%	6%	47%			44%	16%	6%	35%

Table 4.6: Netbid vs. EMSRb Leg Passenger Load by Class Results

Netbid performed poorly in the case above, losing 0.66% in revenue. Compared to the EMSRb vs. EMSRb base case, the airline using Netbid carried significantly more passengers on the high, medium, and low legs. However, by looking at the fare class passenger mix, it was apparent that the airline using Netbid carried a lot more Q class

passengers while it carried fewer M, B, and Y class passengers. Even though the airline using Netbid carried more passengers, its fare class passenger mix was worse than the base case, leading to a loss in revenue. This supported the small network effect theory that the gap between bid-prices were too large with a network of this size, thus accepting on too many lower yield passengers.

Interestingly, in this example case the competitor to the airline using Netbid also suffered in revenue as well. By looking at the leg load results, we can see that because the airline using Netbid let on too many Q passengers, it actually caused the competing airline using EMSRb fare class control to lose some passengers. Most of these passengers were M or Q class passengers, who probably were able to get Q class seats on the Netbid airline instead. In this case, the airline using Netbid not only filled their own airplane with low fare passengers, but they also stole some of the low fare passengers from the competitor. This is one case where both airlines lose because one airline moved to an O-D control method.

The third O-D control RM method in this experiment was DAVN. DAVN performed the best under most cases, and it would be interesting to see why it performed so well. Their passenger load by fare class on three example legs is presented in Table 4.7.

DF	Airline A								Airline B					
	Airline A	Airline B	Leg	Passenger Load by Class					Leg	Passenger Load by Class				
1.0	DAVN	EMSRb	High	Total	Y	B	M	Q	High	Total	Y	B	M	Q
	compared to base				51%	15%	6%	28%			47%	17%	20%	16%
	1.46%	-2.03%	Medium	85.05	37.33	12.39	10.24	25.09	Medium	88.31	36.25	13.17	11.17	27.72
					44%	15%	12%	30%			41%	15%	13%	31%
			Low	72.01	24.28	9.03	4.48	34.22	Low	55.06	23.99	8.55	3.05	19.47
					34%	13%	6%	48%			44%	16%	6%	35%

Table 4.7: DAVN vs. EMSRb Leg Passenger Load by Class Results

The airline using DAVN performed very well, with a 1.46% revenue gain compared to the base case. By looking at the passenger load numbers, it is evident that the airline using DAVN carried fewer passengers on the high and medium legs, while it carried a lot more passengers on the low demand leg. It did carry more Q passengers on the high

demand leg, but it also carried more Y passengers. However, by carrying almost 13.5 extra passengers on the low demand leg, which are the longer haul legs under Network 2, DAVN was able to get a big revenue advantage over the base case.

The airline competing with DAVN suffered a 2.03% loss in revenue. Compared to the base case, it saw a lower yield passenger fare class mix because it carried fewer Y class passengers overall. It also carried much fewer passengers on the low demand legs, which were the longest legs that had the highest revenue value. Basically, compared to the airline using DAVN, the airline using EMSRb fare class control carried slightly more passengers on the high and medium legs, but it carried 17 fewer passengers on the low demand leg. This gave the airline using DAVN a big revenue advantage.

Finally, the last O-D control RM method in this experiment was HBP. HBP was the second-best O-D control RM method, performing almost equally well or slightly behind DAVN in most cases. Its passenger load by fare class on three example legs are presented in Table 4.8.

DF	Airline A								Airline B					
	Airline A	Airline B	Leg	Passenger Load by Class					Leg	Passenger Load by Class				
				Total	Y	B	M	Q		Total	Y	B	M	Q
1.0	HBP	EMSRb	High	90.13	43.23	13.68	8.95	24.27	High	90.42	42.73	15.13	14.59	17.97
	compared to base				48%	15%	10%	27%			47%	17%	16%	20%
	1.32%	-1.78%	Medium	88.03	37.02	12.56	12.03	26.42	Medium	87.57	36.71	12.98	8.83	29.05
					42%	14%	14%	30%			42%	15%	10%	33%
			Low	71.27	24.21	9.08	4.33	33.65	Low	55.74	24.07	8.64	3.07	19.96
					34%	13%	6%	47%			43%	16%	6%	36%

Table 4.8: HBP vs. EMSRb Leg Passenger Load by Class Results

Compared to the base case, the airline using HBP performed 1.32% better. Most of the revenue gain came from the 11 extra passengers the airline using HBP carried on the low-demand, long-haul leg. Compared to DAVN, HBP performed not as well because it had a slightly lower yield passenger fare class mix on the high demand leg, and it had a slightly lower passenger load in Q class on the low-demand leg. Otherwise, its passenger load by class distribution was almost the same as the DAVN case.

As for the airline competing against the airline using HBP, it suffered almost the same fate as the airline competing against DAVN. Compared to the EMSRb base case, it carried about the same number of passengers on the high and medium demand leg, but with a lower yield fare class passenger mix. On the low demand leg, it carried 4.18 fewer passengers than in the base case, mostly in M class. This resulted in a 1.78% loss in revenue.

To sum up, compared to the base case where both airlines use EMSRb, the airlines using O-D control RM methods seem to gain a revenue advantage because they carry more passengers on the long-haul, low-demand leg. They also carry fewer passengers on the high and medium demand legs, but with a better passenger fare class mix. As for the airlines competing against these airlines using O-D control RM methods, they suffered a big loss in revenue because their long-haul passengers were taken away by the competing airline. From the above results we can see how the O-D control RM methods have an advantage over leg based EMSRb fare class control.

After reviewing the leg load results, the next step is to review the path passenger load results. Since we are comparing O-D control methods versus EMSRb fare class control, it is important to look at both the leg passenger loads as well as path passenger loads. Since each O-D market is served by one path from each airline, the path results also represent the number of passengers traveling in a particular market for that airline. For the path loads presented here, the first two are local paths, where the two airlines do not compete with each other. The local results presented here are passenger loads from each airline's hub to a high and a low demand spoke city. The next three sets of results are connecting paths that represent three different markets in which the two airlines compete with each other. The first of the three has paths traversing two low demand legs; the second of the three has paths traversing two high demand legs; finally the last of the three has paths traversing one low demand leg followed by a high demand leg.

Airline A - EMSRb						Airline B - EMSRb					
Passenger Load by Class						Passenger Load by Class					
Path	Total	Y	B	M	Q	Path	Total	Y	B	M	Q
High	39.80	19.49	6.61	9.06	4.64	High	39.89	19.41	6.58	9.13	4.77
Low	14.67	5.35	2.08	0.72	6.52	Low	14.69	5.4	2.11	0.7	6.48
Low-Low	14.31	5.49	1.97	0.76	6.09	Low-Low	14.26	5.5	1.96	0.73	6.07
High-High	16.67	7.78	2.74	3.68	2.47	High-High	16.61	7.79	2.71	3.7	2.41
High-Low	17.27	7.97	2.75	3.67	2.88	High-Low	17.29	7.97	2.76	3.69	2.87

Table 4.9: EMSRb vs. EMSRb Base Case Path Passenger Load by Class Results

In Table 4.9, the base case EMSRb vs. EMSRb passenger load by class results are presented. At a demand factor of 1.0, the local market with a high demand had almost 40 passengers, while the local market with a low demand had only 15 passengers. As for connecting paths, the three example markets chosen for this investigation had a passenger load of 14.3 to 17.3 passengers. It is interesting to see that the path labeled High-Low had a higher passenger load than the path labeled High-High. This is because the results here are path results, not leg results. Even though the two legs a path traverses both have very high total passenger loads, these may be because of a higher local load, or higher loads on other paths that traversed the legs. This EMSRb vs. EMSRb result will be used for comparison to investigate how the O-D control RM methods differed from the base case results.

First, we can look at GVN vs. EMSRb path load results. In Table 4.10 the passenger load by class results are presented for the same selected paths.

Airline A - GVN						Airline B - EMSRb					
Compared to EMSRb : 0.26%						Compared to EMSRb : -1.58%					
Passenger Load by Class						Passenger Load by Class					
Path	Total	Y	B	M	Q	Path	Total	Y	B	M	Q
High	25.06	19.9	3.96	0.34	0.86	High	41.39	19.34	6.73	7.84	7.48
Low	14.37	5.31	2.06	0.69	6.31	Low	14.67	5.44	2.09	0.72	6.42
Low-Low	14.33	5.32	1.98	0.77	6.26	Low-Low	14.15	5.47	1.98	0.71	5.99
High-High	19.62	7.49	2.83	2.69	6.61	High-High	16.31	7.86	2.76	2.44	3.25
High-Low	26.04	7.85	2.99	2.21	12.99	High-Low	14.95	8.06	2.69	1.01	3.19

Table 4.10 – GVN vs. EMSRb Base Case Path Passenger Load by Class Results

Compared to the base case, the airline using GVN had a 0.26% gain in revenue. By looking at the difference in passenger loads per path, we can see that the airline using GVN carried significantly fewer passengers on the high-demand local path, while it carried a lot more passengers on the connecting High-High and High-Low paths. This is

mainly because GVN is greedy, and it prefers connecting passengers with higher fares over local passengers. As a result, the airline using GVN carried fewer local passengers but more connecting passengers. On the path that traverses a high and a low demand leg in particular, this effect is especially strong, leading to 26.04 passengers carried versus the base case with 17.27 passengers carried. On this path with EMSRb fare class control, connecting passengers in M or Q classes get displaced by local Y or B passengers on the high demand leg. However, these connecting passengers, even though they were in a lower fare class, would bring in more total revenue. GVN addresses this by carrying these passengers and displacing local passengers. Therefore, GVN is able to enjoy a small revenue advantage over the airline still using EMSRb.

As for the competitor competing against the airline using GVN, it carried more passengers on the High demand path and fewer passengers on the High-Low demand path. This is obviously because those passengers were carried by the airline using GVN instead. This competing airline actually saw a 1.58% loss in revenue versus the base case.

Moving on to Netbid vs. EMSRb path load results, Netbid performed poorly under most cases, and its path passenger loads show the reason for its poor performance. Table 4.11 shows the passenger load by class results for Netbid vs. EMSRb for the same selected paths.

Airline A - Netbid						Airline B - EMSRb					
Compared to EMSRb : -0.66%						Compared to EMSRb : -1.05%					
Passenger Load by Class						Passenger Load by Class					
Path	Total	Y	B	M	Q	Path	Total	Y	B	M	Q
High	29.67	19.37	4.26	2.4	3.64	High	40.46	19.3	6.75	8.46	5.95
Low	14.55	5.36	2.08	0.69	6.42	Low	14.55	5.36	2.1	0.73	6.36
Low-Low	14.30	5.25	1.96	0.74	6.35	Low-Low	14.38	5.58	1.99	0.74	6.07
High-High	15.52	7.23	2.49	1.03	4.77	High-High	18.79	8.15	3.01	4.46	3.17
High-Low	26.17	7.65	2.91	2.28	13.33	High-Low	14.44	8.1	2.73	1.03	2.58

Table 4.11 – Netbid vs. EMSRb Base Case Path Passenger Load by Class Results

Compared to the EMSRb vs. EMSRb base case, the airline using Netbid carried about 10 fewer passengers on the high demand local path, but it also carried about 9 extra passengers on the high-low demand path. Its passenger fare class mix is slightly worse

than under the EMSRb vs. EMSRb base case as well, with Netbid carrying almost no passengers in M class and many more in Q class. This again suggest that with the big gaps between bid-prices, Netbid allowed too many low fare passengers.

As for the airline competing against the airline using Netbid, it carried slightly more passengers on the High demand local path, but it carried significantly fewer passengers on the High-Low path because these passengers traveled on the Netbid airline in Q class instead. Its fare class passenger mix is also a little bit worse than the base case, and resulted in a much lower revenue.

The third O-D control RM method in this experiment was DAVN, and it was the best performer under most cases. The passenger load by fare class of the airline using DAVN on the five example paths is presented in Table 4.12.

Airline A - DAVN						Airline B - EMSRb					
Compared to EMSRb : 1.46%						Compared to EMSRb : -2.03%					
Passenger Load by Class						Passenger Load by Class					
Path	Total	Y	B	M	Q	Path	Total	Y	B	M	Q
High	27.48	20.37	4.97	0.55	1.59	High	40.13	19.27	6.68	8.47	5.71
Low	14.58	5.39	2.04	0.7	6.45	Low	14.72	5.39	2.11	0.72	6.5
Low-Low	14.43	5.31	1.96	0.78	6.38	Low-Low	14.02	5.39	1.94	0.67	6.02
High-High	13.03	8.03	2.25	0.94	1.81	High-High	19.92	7.7	3.13	5.56	3.53
High-Low	27.17	8.31	2.95	2.02	13.89	High-Low	13.86	7.75	2.69	0.96	2.46

Table 4.12: DAVN vs. EMSRb Base Case Path Passenger Load by Class Results

By comparing the passenger loads on these paths between DAVN and EMSRb, we noticed that the biggest difference in terms of passenger loads is on the local path with high demand. The airline using EMSRb carried 39.80 passengers, while the airline using DAVN carried only 27.48 passengers, most of them in Y class. Furthermore, on the connecting paths that traversed two high demand legs, the DAVN airline also carried fewer passengers, with more of them in Y class. Finally, on the High-Low path, the DAVN airline carried 27.17 passengers, versus the EMSRb airline's 17.27 passengers. As for passenger fare class mix, on four out of five paths described here, the airline using DAVN carried more Y passengers than the airline using EMSRb. Therefore, we can conclude that DAVN performed well by allocating seats to connecting passengers on the higher demand legs instead of to local passengers. With fewer seats for local passengers,

they were able to restrict B, M, and Q class bookings, and carry more Y class passengers. As for those connecting passengers, most of them were on the paths that traverse a high demand leg and a low demand leg. Under EMSRb, these passengers would have been spilled because seats were not available to them on the high demand leg. Instead, by carrying these passengers, DAVN carriers a better passenger fare class mix on the high demand legs, while retaining these connecting passengers who may be paying a high connecting fare.

As for the airline competing against the DAVN airline, it ended up carrying fewer passengers in Y class and more passengers in lower classes in most paths. It also carried fewer passengers on the connecting path that traversed a high demand and a low demand leg. With a lower passenger load as well as a poorer passenger fare class mix, their revenue result compared to the base case was much worse.

Finally the last O-D control RM method is HBP. HBP performed almost as well as DAVN in most cases. Its passenger load by fare class is presented in Table 4.13 for the example paths.

Airline A - HBP						Airline B - EMSRb					
Compared to EMSRb : 1.32%						Compared to EMSRb : -1.78%					
Passenger Load by Class						Passenger Load by Class					
Path	Total	Y	B	M	Q	Path	Total	Y	B	M	Q
High	27.32	19.48	5.39	0.51	1.94	High	40.74	19.28	6.72	7.91	6.83
Low	14.62	5.38	2.08	0.72	6.44	Low	14.69	5.45	2.13	0.7	6.41
Low-Low	14.45	5.38	1.97	0.76	6.34	Low-Low	14.14	5.53	1.95	0.73	5.93
High-High	16.88	7.86	2.62	3.31	3.09	High-High	17.82	7.77	2.85	3.1	4.1
High-Low	26.69	8	2.97	1.95	13.77	High-Low	14.37	7.82	2.7	0.97	2.88

Table 4.13: HBP vs. EMSRb Base Case Path Passenger Load by Class Results

By comparing the passenger load by class results of the airline using HBP to the airline using DAVN, it is obvious that under both methods, fewer local passengers were carried on the high demand path, while more connecting passengers were carried on the paths that traversed a high and a low demand leg. This was why both HBP and DAVN had such a large revenue advantage over the base case airline using EMSRb. HBP performed slightly worse than DAVN mainly because its passenger fare class mix was not as good. It carried a smaller number of Y passengers and slightly more passengers in other fare

classes. However, it still performed 1.32% better than the base case airline using EMSRb fare class control.

As for the airline competing against the airline using HBP, it suffered almost the identical fate as the airline competing against the airline using DAVN. More local passengers were carried on the high demand path, while fewer passengers were carried on the connecting paths that traversed through a high and a low demand leg. That, coupled by a lower yield passenger fare class mix, resulted in 1.78% loss in revenue compared to the base case airline using EMSRb fare class control.

To sum up, compared to the base case where both airlines use EMSRb, the airlines using O-D control RM methods gain a revenue advantage because they carried more higher fare paying passengers. They carried a better fare class passenger mix on most legs, resulting in more Y class passengers and higher revenue. Also, on the high demand “bottleneck” legs, they carried more passengers on connecting paths in place of passengers on local paths, thus bringing in even higher revenue. As for the airlines competing against airlines using O-D control methods, on the high demand legs they ended up carrying more passengers on local paths instead of passengers on connecting paths. Since local passengers generally bring in less network revenue than connecting passengers, and since some Y class passengers were stolen by the airline using O-D control, the competing airline suffered in revenue when compared to the base case. From the results here, again we can see the advantages of using O-D control Revenue Management methods over EMSRb fare class control.

4.1.4 Summary

In this section the results from the experiments on Origin-Destination control Revenue Management methods will be summarized. The relative performance of the four O-D control RM methods under the three different Networks will be discussed, followed by the explanations of the better performance from the O-D control methods.




Figure 4.1: The performance of O-D control RM methods under Network 1

Figure 4.1 shows the relative performances of the O-D methods under Network 1, which consists of a co-located hub and four spoke cities. Under all demand factors, DAVN performed the best, followed by HBP. Their revenue performance increased relative to the base case as the demand factor increased. GVN performed well under a demand factor of 0.8, but its performance progressively worsened as demand increased. Its revenue performance was also behind DAVN and HBP. Finally, Netbid was the worst performer under all demand factors under Network 1.

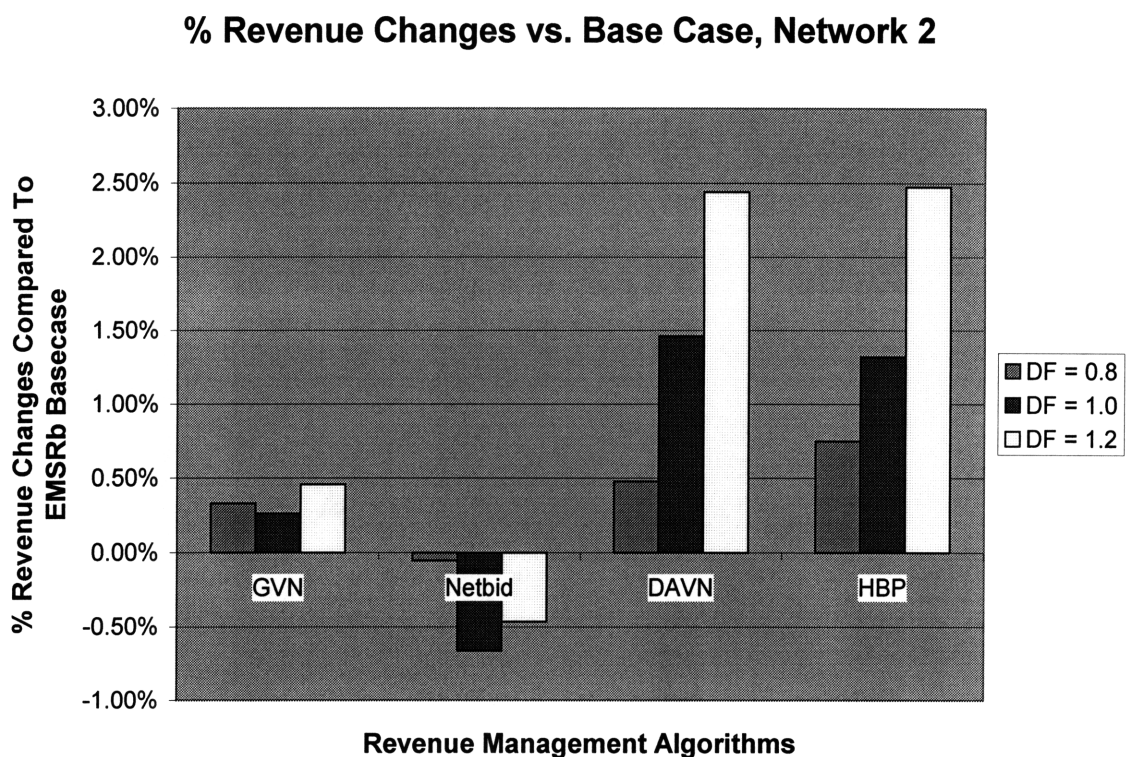


Figure 4.2: The performance of O-D control RM methods under Network 2

Figure 4.2 shows the relative performances of the O-D methods under Network 2, which consists of one hub for each airline, plus four spoke cities. Again, DAVN and HBP performed the best, but this time HBP performed slightly better than DAVN. Both methods also performed better when the demand factor increased. GVN’s performance again lagged behind DAVN and HBP, but did not result in revenue losses. Finally, Netbid again was the worst performer of the group, resulting in a decrease in revenue when compared to the base case EMSRb fare class control method.

% Revenue Changes vs. Base Case, Network 3




Figure 4.3: The performance of O-D control RM methods under Network 3

Figure 4.3 shows the relative performance of O-D methods under Network 3, which consisted of one hub for each airline, plus six spoke cities. DAVN performed the best, followed by HBP. They both performed better as the demand increased. GVN again showed positive results, but did not perform as well as the first two O-D control methods. Finally, Netbid showed positive results, but its results were still the worst of the four O-D methods.

One thing to note about GVN is that under Network 1, its performance worsened as the demand factor increased. This is because GVN gives preference to connecting passengers over local passengers, and when the demand is high, it might be displacing two high-fare local passengers by carrying one connecting passenger. The two high-fare local passengers then traveled on the competing airline instead, and GVN revenues suffered as a result. Under Networks 2 and 3, since the two airlines have their own hubs, the displaced local passengers could not travel on the other airline, and in fact might have “sold up” to a higher fare on the same airline, so this effect was not as pronounced as

under Network 1. For a further discussion of Network 1 vs. Network 2, please see Section 4.4.

Another thing to note is that for Netbid, it performed badly under all three networks. One reason was because the pick-up forecasting and booking curve detruncation used in the experiments just did not work as well with Netbid. Another reason was because with such a small network, the bid prices obtained from the deterministic network solution were too far apart as the flights booked up, and Netbid ended up letting too many low-fare paying passengers fly. This “small network effect” caused Netbid to perform worse in Networks 1 and 2, which have only 4 spoke cities. Netbid actually achieved a positive revenue result under Network 3, which had 6 spoke cities.

To determine why the O-D based methods performed better than the EMSRb fare class control, we examined both leg and path load differences. In terms of leg loads, the airline using O-D control RM methods performed better because it carried more passengers on the long-haul, low-demand legs. It also carried a better passenger fare class mix on the high and medium demand legs. As for the competing airline, its long-haul passengers were taken away by the airline using O-D control methods, leading to a loss in revenue.

By looking at the results on a path basis, we can see that the airline using O-D control methods performed better because on the bottleneck high demand legs, it carried more connecting passengers instead of local passengers. These connecting passengers travel on paths that traverse a high demand and a low demand leg. They pay a high total fare, but were being displaced under EMSRb fare class control. The lower number of local passengers also led to a better fare class passenger mix. All these contribute to the success of O-D control Revenue Management methods.

In this section, comparisons between GVN, Netbid, DAVN, and HBP were made to EMSRb under different networks and demand factors. The reason why these O-D control methods performed better were also discussed. In short, under the conditions

tested here, DAVN was the best performing O-D method, followed closely by HBP. GVN was the next best, but it did not perform well under all conditions. Finally, Netbid performed the worst under all conditions, leading to a loss in revenue in most cases.

4.2 Impacts of Passenger Choice Assumptions

4.2.1 Motivation

PODS simulates how a passenger chooses which path and fare class to travel on through the Passenger Choice Model discussed earlier in Section 2.2. Since the passenger choice model has a great impact on the simulation results, it is important to investigate the effects of different passenger choice model assumptions. For this investigation, first, different levels of demand variation as well as passengers' ability to choose an alternative path was investigated. The three levels of passenger choice and demand variation are called Sequences 0, 2, and 3, as discussed in Section 3.4. The experiment results for these three sequences will be presented in Section 4.2.3. Finally, a new passenger choice formulation, as described in Section 2.4.8, is investigated. The results will be presented in Section 4.2.4.

4.2.2 Simulation Setup

Network 2 with 4 spoke cities and separate hubs for each airline was used for both passenger choice experiments. For the Sequence 0, 2, and 3 experiments, Demand Factors of 0.8, 1.0, and 1.2 were used. This was done to see the effect increasing demand has on the revenue results under the three different sequence assumptions. For the second experiment, the PODS7 maximum willingness-to-pay passenger choice formulation was used for comparisons between the old, PODS6 linear formulation and the new, PODS7 exponential formulation. Both experiments were run with the first airline using different O-D control RM methods, and the second airline using EMSRb fare class control. The comparison percentage numbers are again expressed relative to the base case, which was again EMSRb vs. EMSRb. As for the PODS7 passenger choice

model representation experiments, further simulations were run with both airlines using different combinations of O-D control RM methods.

4.2.3 Sequence 0 and Sequence 2 vs. Sequence 3

This section presents the results from the experiments on Sequence 0, Sequence 2, and Sequence 3. Section 3.4 already discussed the three sequences, so here only a summary is given:

Under Sequence 0, path/fare class demand is assumed to be known and independent of availability. Passengers make their preferred travel choice based on the passenger choice model, and this choice ends up being their only choice. If their first choice itinerary is not available, they do not travel and get spilled. Furthermore, the system, market, and passenger type K- factors are set to zero, creating total ODF independence. Finally, the Demand Z-factors are set to 1.0. This combination of assumptions result in independent ODF demands, and passengers do not sell-up and only fly on their first choice itinerary.

Sequence 2 is very similar to Sequence 0, except that now increased demand variation is introduced. Sequence 2 still assumes that path/fare class demand is known and independent of availability, and passengers only travel on their first choice. However, ODF demands are correlated with each other and are more variable than in Sequence 0.

Finally, Sequence 3 introduces full path/fare class choice dependent on flight availability. Passengers can now choose other itineraries to travel on, if their first choice is not available. Again, ODF demands are correlated with each other and are more variable than in Sequence 0.

With the three sequences explained, now we are ready to move onto the experiment results. Table 4.14 shows the results for Sequences 0, 2, and 3 under a Demand Factor of 0.8.

			Sequence 0		Sequence 2		Sequence 3	
			% increase		% increase		% increase	
			compared to		compared to		compared to	
YM methods			EMSRb	EMSRb	EMSRb	EMSRb	EMSRb	EMSRb
DF	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B
0.8	EMSRb	EMSRb						
0.8	GVN	EMSRb	0.04%		0.51%		0.33%	-0.30%
0.8	Netbid	EMSRb	0.07%		-0.10%		-0.05%	0.22%
0.8	DAVN	EMSRb	0.04%		0.55%		0.48%	-0.55%
0.8	HBP	EMSRb	0.04%		0.53%		0.75%	-0.30%

Table 4.14: Sequences 0 vs. 2 vs. 3, DF = 0.8

At this low demand factor, the changes in revenue compared to the base case, when both airlines use leg based EMSRb control, was not very big. Under Sequence 0 with no passenger choice and independent ODF demand, the improvements from O-D based methods were only 0.04% to 0.07%. Under Sequence 2, by introducing ODF demand correlation and greater variance but still no passenger choice, the O-D based methods saw a 0.51% to 0.55% gain in revenue, except Netbid which again performed poorly. By introducing full passenger choice and ODF demand correlation under Sequence 3, we see that the impacts of O-D control methods are somewhat smaller under 3 of the 4 methods. This is due to the fact that with passenger choice, there are now some interactions between the loads and revenues of the two airlines, thus leading to a slight reduction in improvements in revenue when compared to Sequence 2.

As for the airline competing against the airline with an O-D control method, under Sequences 0 and 2 its revenues stay the same as the base case. This was because with no passenger choice, passengers cannot switch airlines, so the two airlines act as independent competitors, and their revenue results stay the same. Under Sequence 3, when one competitor used an O-D control RM method other than Netbid, its revenue suffered while the competitor gained. Letting the passengers choose between airlines under Sequence 3, the airline still using EMSRb loses some valuable passengers to the airline using O-D control, thus leading to a loss in revenue.

The next set of results is for Demand Factor = 1.0, presented in Table 4.15.

			Sequence 0		Sequence 2		Sequence 3	
			% increase		% increase		% increase	
			compared to		compared to		compared to	
	YM methods		EMSRb	EMSRb	EMSRb	EMSRb	EMSRb	EMSRb
DF	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B
1.0	EMSRb	EMSRb						
1.0	GVN	EMSRb	0.84%		0.89%		0.26%	-1.58%
1.0	Netbid	EMSRb	0.64%		-0.28%		-0.66%	-1.05%
1.0	DAVN	EMSRb	0.93%		1.31%		1.46%	-2.03%
1.0	HBP	EMSRb	0.90%		1.17%		1.32%	-1.78%

Table 4.15: Sequences 0 vs. 2 vs. 3, DF = 1.0

At DF = 1.0, the effects of the O-D control RM methods were much bigger. Under Sequence 0, with no passenger choice and independent ODF demand, the revenue differences compared to base case ranged from 0.64% to 0.93%. Under Sequence 2 with no passenger choice but with ODF demand correlation and increased variance, the revenue difference when compared to base case increased slightly except for Netbid. Netbid saw a positive revenue gain under Sequence 0, but saw revenues decrease as passenger choice and greater demand variability were introduced. Finally, for Sequence 3 with full passenger choice and ODF demand correlation, DAVN and HBP saw an even bigger percentage of revenue gain. GVN and Netbid on the other hand, saw a large percentage revenue decrease. The results between Sequences 0 and 2 tell us that the relative performance of O-D control methods compared to EMSRb increase when ODF demand correlation and demand variation are introduced into the simulation. The results between Sequences 2 and 3 tells us that by introducing passenger choice and allowing the two airlines to compete with each other, the percentage change in revenue results are substantially different for different O-D methods.

As for the airline competing against the airline using O-D control RM methods, under Sequence 3 it saw a large decrease in revenue when compared to the base case. This tells us that with interaction and competition between the two airlines, the airline using O-D control RM methods can make the airline using fare class control suffer in terms of revenue. This again is another example of the “zero-sum” effect where the revenue gain of the airline using O-D control methods comes partly at the expense of the competing airline using EMSRb fare class control.

Finally, moving onto $DF = 1.2$, the results are presented in Table 4.16.

	YM methods		Sequence 0		Sequence 2		Sequence 3	
			% increase		% increase		% increase	
			compared to		compared to		compared to	
DF	Airline A	Airline B	EMSRb	EMSRb	EMSRb	EMSRb	EMSRb	EMSRb
1.2	EMSRb	EMSRb						
1.2	GVN	EMSRb	0.85%		0.64%		0.46%	-0.99%
1.2	Netbid	EMSRb	1.34%		-0.30%		-0.47%	-0.58%
1.2	DAVN	EMSRb	1.50%		1.77%		2.44%	-1.67%
1.2	HBP	EMSRb	1.55%		1.57%		2.47%	-1.56%

Table 4.16: Sequences 0 vs. 2 vs. 3, $DF = 1.2$

Under $DF = 1.2$, again we see the same results as under $DF = 0.8$ and 1.0 , except with a bigger magnitude of percentage change in revenue compared to the base case. Netbid still saw positive revenue results when compared to base case under Sequence 0, but once passenger choice and demand correlation were introduced, Netbid results became negative. GVN results also suffered when we moved from Sequence 0 to 2 and 3. Finally for DAVN and HBP, as passenger choice and demand correlation were introduced, their performance compared to the base case improved as well. For the competing airline under Sequence 3, again its revenue results suffered. This again showed evidence of the “zero-sum” effect once the demand of the two airlines interact with each other.

To sum up, as demand factors increased, the percentage change in revenue of O-D control RM methods compared to the base case increased as well. As for individual RM algorithms, GVN results were inconsistent when we moved from Sequences 0 to 2 and 3. Under different demand factors, sometimes GVN results improved, while for others they decreased. Netbid results were always positive under Sequence 0, but once demand correlation and passenger choice were introduced, Netbid revenue results suffered and went into the negative territory when compared to the base case. Finally, for DAVN and HBP, as we introduced demand correlation and passenger choice, their performance relative to EMSRb base case improved as well. As for the competing EMSRb airline under Sequence 3, with passenger choice and demand correlation, its revenue results suffer because the airline using O-D control RM methods were able to control seat

availability better so that they captured a better passenger fare class mix with higher revenue.

4.2.4 PODS7 Representation of Passenger Choice

The demand under PODS6 is controlled by a maximum willingness-to-pay parameter, which determines the maximum fare price that 50% of the passengers are willing to pay. This representation tended to result in very high sell-up rates when the lowest fare class is unavailable. Furthermore, if the fares are dropped, the demand stays the same unless the PODS analyst manually changes it. As a result, we are unable to simulate the competitive situation when one airline drops its fare to stimulate demand.

To overcome the high sell-up rate as well as PODS' inability to simulate demand stimulation, a new passenger choice maximum willingness-to-pay formulation was designed and implemented in PODS version 7. All experiments discussed in this thesis except the ones discussed in this section were run under PODS version 6, with the previous maximum willingness-to-pay formulation. The PODS6 vs. PODS7 experiments presented here is used to compare the effect of the two different maximum willingness-to-pay passenger choice model formulations under the same network.

The PODS7 maximum willingness-to-pay formulation was already discussed in Section 2.4.8, but is summarized briefly here again.

For any given market and passenger type, the new maximum willingness-to-pay formulation is:

$$P(\text{xpay}_{m,t} \geq f) = \min \left[1., e^{-\left[\frac{.6931(f - \text{fb}_{m,t})}{(\text{em}_{m,t} - 1)\text{fb}_{m,t}} \right]} \right]$$

where:

- $\text{xpay}_{m,t}$ = max willing-to-pay, market m, pax type t
- $\text{fb}_{m,t}$ = input base fare, market m, pax type t
- f = fare
- $\text{em}_{m,t}$ = input multiple of base fare (emult) at which the probability a random pax of type t in market m will travel = 0.5

Note that since the demand can be more than 100% of the demand at the *basefare*, when computing the effective demand, this function can be greater than 1, so we can ignore the min[1] part of the above formula. One other thing to note about this new formulation is that it is a downward sloping exponential curve, which is the same shape as a demand curve. This new formulation should represent airline passenger demand curve more closely.

Now, the effective demand is just the above formula multiplied by the demand at *basefare*, so

$$\text{Effective Demand} = \text{Input Demand @ basefare} \times e^{-\left[\frac{0.6931(f - \text{fb}_{m,t})}{(\text{em}_{m,t} - 1)\text{fb}_{m,t}} \right]}$$

The elasticity of demand with respect to price is then:

$$(\varepsilon_{m,t})_{D,P} = -\left(\frac{f}{fb_{m,t}}\right) \times \left(\frac{0.6391}{em_{m,t} - 1}\right)$$

where:

$(\varepsilon_{m,t})_{D,P}$ = elasticity of demand with respect to price, for market m , pax type t .

In the experiments discussed here, the following parameters were used:

For an example market of 1000 miles, the fares for Y, B, M, and Q classes are \$800, \$400, \$300, and \$200, respectively. For business travelers, their *basefare* = 500, and their *emult* = 3. For Leisure travelers, their *basefare* = 200, and their *emult* = 1.2. This means that 50% of Business Travelers are willing to pay 3 times the *basefare* to travel, and 50% of leisure travelers are willing to pay 1.2 times the *basefare* to travel. This results in the following elasticities with respect to price, at different fare levels:

Fare	1200	1000	800	600	400	300	200
Elasticity - Business Traveler	-0.83172	-0.6931	-0.55448	-0.41586	-0.27724	-0.20793	-0.13862
Elasticity - Leisure Traveler	-20.793	-17.3275	-13.862	-10.3965	-6.931	-5.19825	-3.4655

Table 4.17: Demand Elasticity with Respect to Price for Business and Leisure Travelers

Throughout the price ranges of the markets in the networks covered by this thesis, the demand elasticity for the Business Travelers is always inelastic, and the demand elasticity for the Leisure Travelers is always elastic. The values of the demand elasticities are reasonable as they correctly represent the characteristics of each type of passengers. Also, as the fares go up, the demand elasticity becomes more elastic for both types of passengers. This is again reasonable because higher fares result in more elastic demand.

Next, the results from the PODS7 new passenger demand formulation are compared to PODS6 results, and presented here. We'll start with $DF = 0.8$:

				PODS6		PODS7	
				% increase		% increase	
				compared to		compared to	
		YM Method		EMSRb	EMSRb	EMSRb	EMSRb
DF	Seq	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B
0.8	3	GVN	EMSRb	0.77%	-0.70%	0.69%	-0.58%
0.8	3	DAVN	EMSRb	1.14%	-0.66%	0.82%	-0.64%
0.8	3	Netbid	EMSRb	0.01%	-0.54%	0.14%	-0.28%
0.8	3	HBP	EMSRb	0.73%	-0.83%	0.63%	-0.53%

Table 4.18: PODS6 vs. PODS7 Revenue Results, DF = 0.8

From Table 4.18, first we notice that with the new formulation under PODS7, the relative rankings of O-D control RM methods have not changed. DAVN still performed the best, followed by GVN and HBP. Netbid still performed the worst. Figure 4.4 also illustrates this trend. Furthermore, it is evident from Figure 4.4 that the revenue performance for O-D control RM methods relative to leg based EMSRb under PODS7 was greater. This is because EMSRb doesn't benefit as much from sell-up, meaning relative gains of O-D control methods are greater under PODS6.

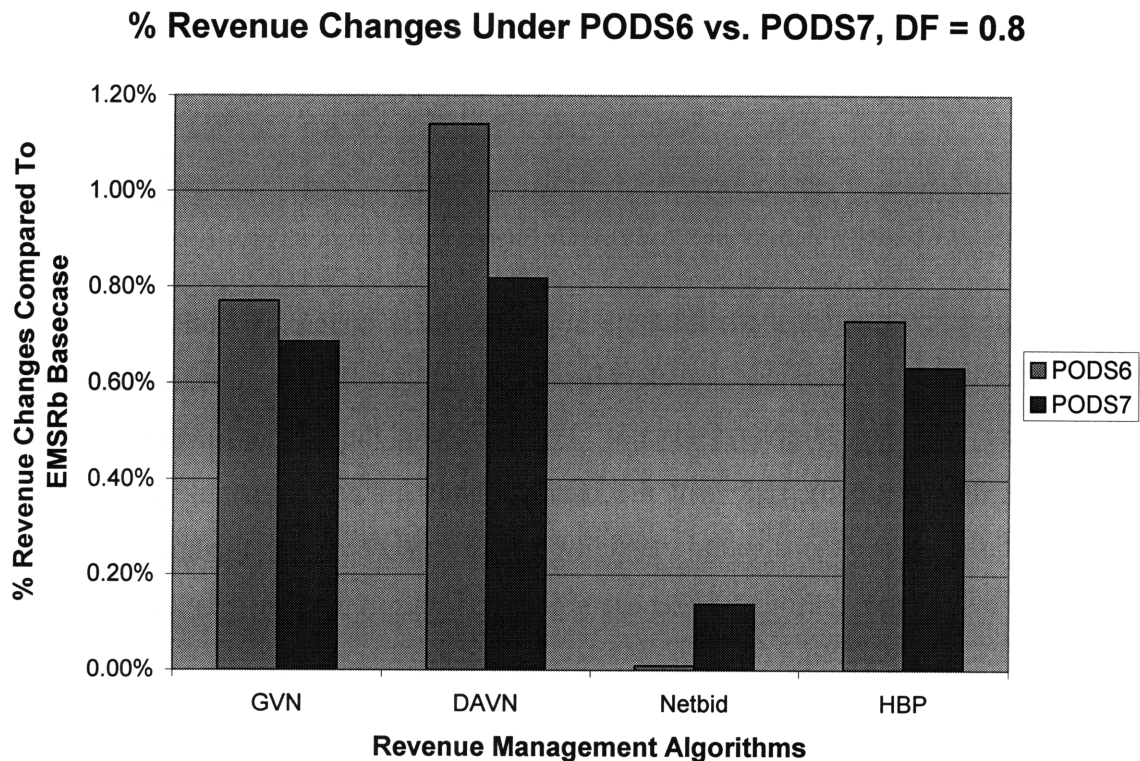


Figure 4.4: % Revenue Changes Compared to EMSRb Under PODS6 vs. PODS7, DF = 0.8

This same phenomenon can again be seen in Table 4.19 where both airlines used O-D control RM methods at DF = 0.8.

				PODS6		PODS7	
				% increase		% increase	
				compared to		compared to	
				YM Method		EMSRb	
DF	Seq	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B
0.8	3	GVN	GVN	-0.07%	-0.11%	0.09%	0.06%
0.8	3	GVN	DAVN	-0.38%	-0.14%	0.04%	0.17%
0.8	3	GVN	Netbid	0.33%	-0.57%	0.36%	-0.44%
0.8	3	GVN	HBP	-0.08%	0.12%	0.11%	0.01%
0.8	3	DAVN	DAVN	0.23%	-0.01%	0.14%	0.12%
0.8	3	DAVN	Netbid	0.27%	-0.65%	0.48%	-0.49%
0.8	3	DAVN	HBP	-0.03%	-0.32%	0.22%	-0.03%
0.8	3	Netbid	Netbid	-0.40%	-0.41%	-0.15%	-0.17%
0.8	3	Netbid	HBP	-0.70%	10.00%	-0.34%	0.15%
0.8	3	HBP	HBP	-0.29%	-0.09%	0.08%	0.04%

Table 4.19: PODS6 vs. PODS7 Revenue Results, O-D RM Methods Competing Head to Head, DF = 0.8

Again, PODS7 gave similar results as PODS6, but with better revenue performance for O-D control RM methods relative to EMSRb.

Moving on to a demand factor of 1.0, again we see the same trend in the results. The relative rankings of RM algorithms stayed the same under both PODS versions. Again DAVN performed the best, followed by HBP, GVN, and finally Netbid. Similarly, the revenue improvements were much greater under PODS7, especially for GVN.

				PODS6		PODS7	
				% increase		% increase	
				compared to		compared to	
				YM Method		EMSRb	
DF	Seq	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B
1.0	3	GVN	EMSRb	0.56%	-1.91%	1.58%	-0.95%
1.0	3	DAVN	EMSRb	2.09%	-1.92%	2.33%	-1.20%
1.0	3	Netbid	EMSRb	0.52%	-1.38%	0.84%	-0.37%
1.0	3	HBP	EMSRb	1.77%	-1.69%	1.97%	-1.00%

Table 4.20: PODS6 vs. PODS7 Revenue Results, DF = 1.0

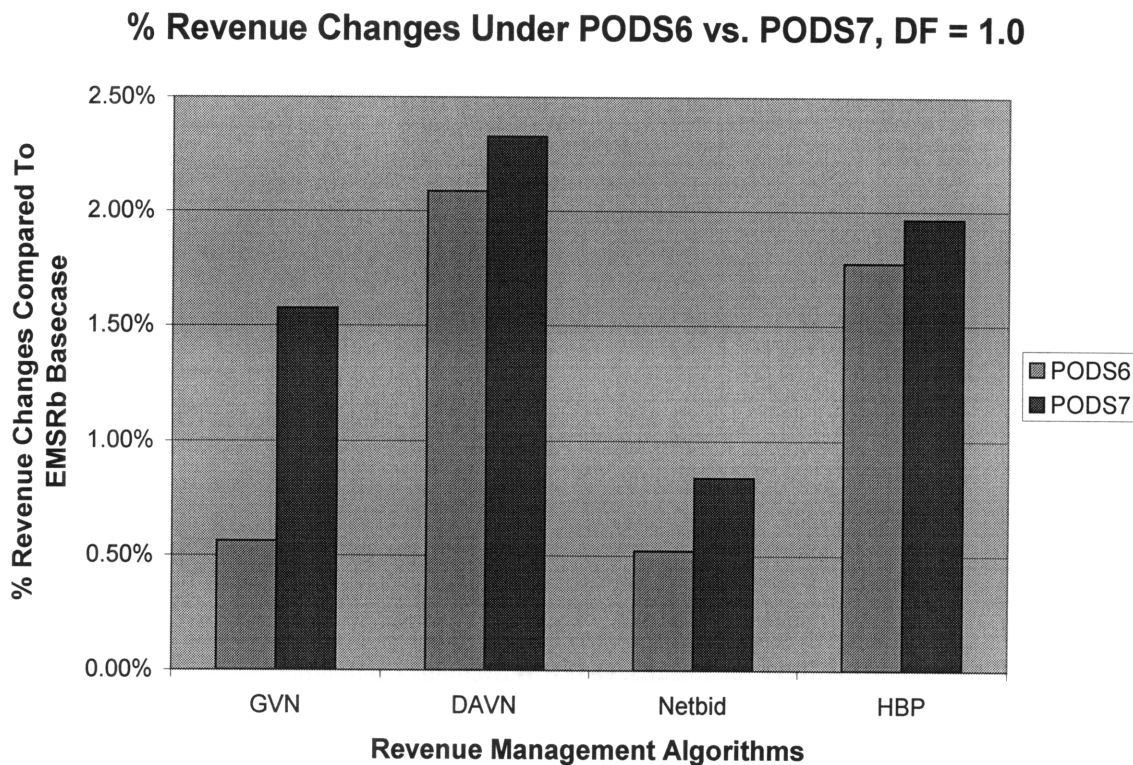


Figure 4.5: % Revenue Changes Compared to EMSRb Under PODS6 vs. PODS7, DF = 1.0

As for head to head results, interestingly under PODS6 when both airlines used O-D control RM methods, they both suffered a revenue loss when compared to EMSRb. Under PODS7 when both airlines used O-D control RM methods, they both see a revenue gain except for Netbid, which still performed poorly. This revenue gain is again due to the lower level of sell-up under PODS7, and since EMSRb benefits the least from sell-up, it made the base case EMSRb airline perform better relative to other methods under PODS7, reducing the revenue differential.

				PODS6		PODS7	
				% increase		% increase	
				compared to		compared to	
		YM Method		EMSRb	EMSRb	EMSRb	EMSRb
DF	Seq	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B
1.0	3	GVN	GVN	-0.58%	-1.01%	0.50%	0.45%
1.0	3	GVN	DAVN	-0.63%	0.26%	0.25%	0.94%
1.0	3	GVN	Netbid	-0.11%	-1.32%	1.04%	-0.48%
1.0	3	GVN	HBP	-0.79%	-0.77%	0.43%	0.53%
1.0	3	DAVN	DAVN	-0.10%	-0.32%	0.80%	0.78%
1.0	3	DAVN	Netbid	0.70%	-1.54%	1.68%	-0.60%
1.0	3	DAVN	HBP	0.24%	-0.59%	0.98%	0.36%
1.0	3	Netbid	Netbid	-0.49%	-0.85%	0.15%	0.10%
1.0	3	Netbid	HBP	-0.82%	-0.96%	-0.19%	0.37%
1.0	3	HBP	HBP	-0.33%	-0.47%	0.60%	0.56%

Table 4.21: PODS6 vs. PODS7 Revenue Results, O-D RM Methods Competing Head to Head, DF = 1.0

Finally, moving onto DF = 1.2 results, Table 4.22 and Figure 4.6 show the revenue change when one airline used O-D control RM methods while the second airline still used EMSRb.

				PODS6		PODS7	
				% increase		% increase	
				compared to		compared to	
		YM Method		EMSRb	EMSRb	EMSRb	EMSRb
DF	Seq	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B
1.2	3	GVN	EMSRb	1.23%	-1.33%	1.35%	-0.64%
1.2	3	DAVN	EMSRb	2.60%	-1.63%	2.71%	-1.03%
1.2	3	Netbid	EMSRb	1.08%	-1.06%	1.06%	0.02%
1.2	3	HBP	EMSRb	2.41%	-1.51%	2.51%	-1.03%

Table 4.22: PODS6 vs. PODS7 Revenue Results, DF = 1.2

The relative rankings of RM algorithms stayed the same under both PODS versions. Again, three out of four RM algorithms saw a bigger revenue increase under PODS7, with Netbid having about the same performance under both PODS versions.

% Revenue Changes Under PODS6 vs. PODS7, DF = 1.2

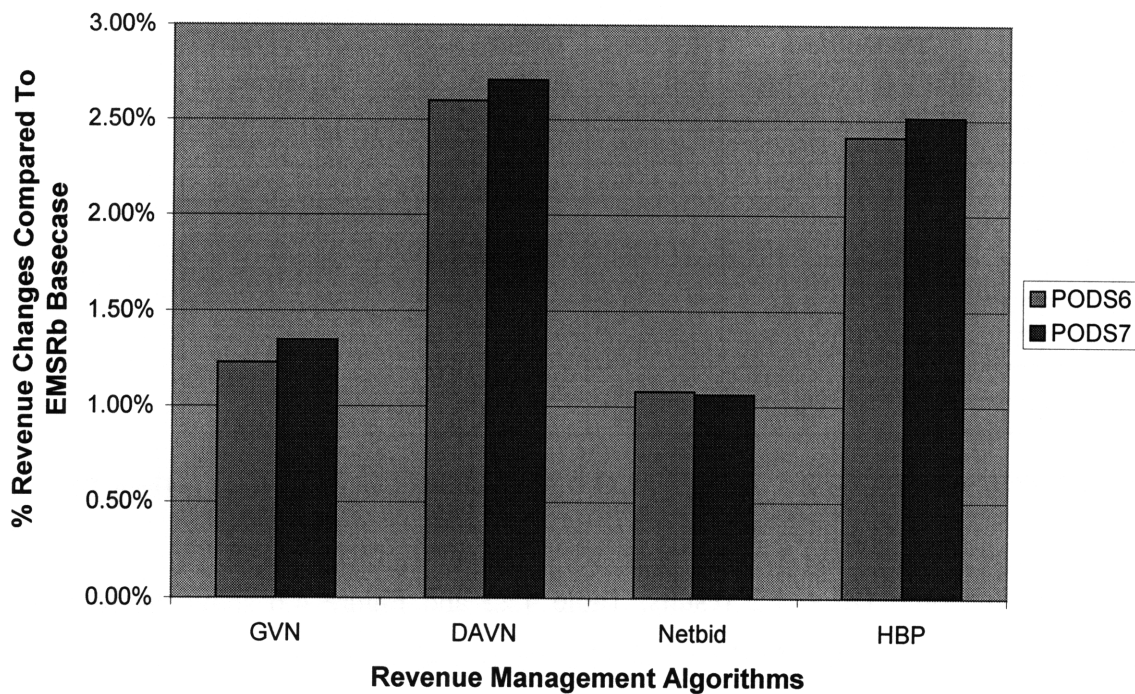


Figure 4.6: % Revenue Changes Compared to EMSRb Under PODS6 vs. PODS7, DF = 1.2

Finally, looking at O-D RM algorithms competing Head to Head at DF = 1.2 in Table 4.23, we again see that all O-D algorithms performed better, or had more positive results, under PODS7. This again was due to the lower level of sell-up and different willingness-to-pay levels for both business and leisure travelers. However, the relative rankings of RM algorithms still stayed the same, with DAVN performing the best, followed by HBP and GVN, and finally Netbid. When both airlines used Netbid, they both saw a revenue improvement versus the leg based EMSRb algorithm. However, when one airline used Netbid against other O-D algorithms, its revenue results suffered.

				PODS6		PODS7	
				% increase		% increase	
				compared to		compared to	
		YM Method		EMSRb	EMSRb	EMSRb	EMSRb
DF	Seq	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B
1.2	3	GVN	GVN	-0.27%	-0.31%	0.42%	0.44%
1.2	3	GVN	DAVN	-0.43%	1.04%	0.02%	1.62%
1.2	3	GVN	Netbid	0.21%	-0.79%	1.00%	-0.19%
1.2	3	GVN	HBP	-0.23%	0.95%	0.02%	1.32%
1.2	3	DAVN	DAVN	1.00%	0.95%	1.40%	1.45%
1.2	3	DAVN	Netbid	1.30%	-0.72%	2.32%	-0.19%
1.2	3	DAVN	HBP	1.10%	0.82%	1.27%	0.99%
1.2	3	Netbid	Netbid	-0.13%	-0.03%	0.50%	0.58%
1.2	3	Netbid	HBP	-0.53%	-0.28%	-0.05%	0.74%
1.2	3	HBP	HBP	0.56%	0.57%	0.97%	1.01%

Table 4.23: PODS6 vs. PODS7 Revenue Results, O-D RM Methods Competing Head to Head, DF = 1.2

Finally, for passenger loads and class distribution, please see Appendix C for a detailed printout. We observed that the total passenger loads were higher under PODS6. Furthermore, due to the lower *emult* for leisure travelers under PODS7, we saw many more Q passengers relative to other fare classes. This is because the next higher fare class, M class, has a fare of 1.5 times the *basefare*. With an *emult* of only 1.2, not too many of the leisure travelers were willing to buy up to M class. As for Y class bookings, we also saw that Y class bookings were much lower, while B class bookings increased slightly. This was because of less sell-up from B class to Y class.

In conclusion, in this section the experiment results for a new maximum willingness-to-pay formulation was presented. With this new passenger choice model, the results were consistent in relative RM algorithm rankings with previous results from the old passenger choice model. We also observed that % increase in revenue compared to EMSRb vs. EMSRb base case was higher for both airlines under all O-D algorithms. The total passenger loads were also slightly lower under PODS7, with slightly higher loads in Q and B classes, but much lower loads in M and Y classes. With this lower level of sell-up under PODS7, we saw lower passenger loads and more spill, as well as a lower absolute revenue. All these made the impact of O-D algorithms greater under the PODS7

formulation, as the revenue gains relative to EMSRb (which no longer benefited from strong sell-up) increased.

4.3 Competing Airlines with Co-located Hub vs. Separate Hubs

4.3.1 Motivation

Network 1 consists of one common hub for both airlines, and airlines compete for local traffic from their hubs to each of the spoke cities. This simulates a situation such as Chicago O'Hare, where both American Airlines and United Airlines have their hubs. However, since most of the domestic US airlines do not co-locate their hub in the same city, it would be interesting to see the effect of separating the hubs to create a separate hub for each airline. Thus, Network 2 was created to simulate the more common situation at most US hubs. It consists again of 4 cities, but this time the hubs do not co-locate in the same city. Both airlines serve the 4 spoke cities, but do not operate to the other's hub city. They compete with each other for connecting traffic, but each has a monopoly on traffic in and out of their respective hubs.

It is of interest to find out the differences in O-D revenue management algorithm performance under the two different network scenarios, and this section discusses the results from this experiment.

4.3.2 Simulation Setup

For this hub location experiment, Networks 1 and 2 were used under PODS6. Please see Figure 4.7 for the topology of Network 1, and Figure 4.8 for the topology of Network 2. Sequences 0, 2, and 3 were all used in the experiment so that the effects of passenger choice and demand variation could be examined. Finally, Demand Factors of 0.8, 1.0,

and 1.2 were used to see the effect of increasing passenger demand on the performance of O-D control RM methods on the two networks.




Figure 4.7: Topology of Network 1




Figure 4.8: Topology of Network 2

4.3.3 Comparison Between Co-located Hubs and Separate Hubs

Table 4.24 shows the results of Network 1 vs. Network 2 under Demand Factor of 0.8, Sequences 0 and 2.

				Network 1		Network 2	
				% increase		% increase	
				compared to		compared to	
				EMSRb	EMSRb	EMSRb	EMSRb
DF	Seq.	YM methods		Airline A	Airline B	Airline A	Airline B
0.8	0	EMSRb	EMSRb				
0.8	0	GVN	EMSRb	0.06%		0.04%	
0.8	0	Netbid	EMSRb	0.07%		0.07%	
0.8	0	DAVN	EMSRb	0.06%		0.04%	
0.8	0	HBP	EMSRb	0.06%		0.04%	
0.8	2	EMSRb	EMSRb				
0.8	2	GVN	EMSRb	0.51%		0.51%	
0.8	2	Netbid	EMSRb	-0.09%		-0.10%	
0.8	2	DAVN	EMSRb	0.56%		0.55%	
0.8	2	HBP	EMSRb	0.53%		0.53%	

Table 4.24: Network 1 vs. Network 2, DF = 0.8, Sequences 0 and 2

Demand Factor of 0.8 means the average demand over all of the legs on the network is 80% of the capacity. At this demand factor, because seat availability is not as constrained, the effects of different algorithms were not as great. Under Sequence 0 and 2 where the assumption was no passenger choice, the different O-D based algorithms should perform the same under Networks 1 and 2. Because passengers did not have a choice in the simulation, there can be no spill or sell-up to the other airline on the connecting itineraries whether the hubs were co-located or not. Local passengers also could not choose the competitor because no passenger choice was allowed. Both airlines operated independently of each other under the Sequence 0 and 2 assumptions, thus making no difference in RM method performance. This is confirmed by the Sequence 0 and 2, Network 1 vs. Network 2 results, where the O-D based RM methods performed equally well under both networks.

		Network 1				Network 2	
		% increase				% increase	
		compared to				compared to	
		YM methods		EMSRb	EMSRb	EMSRb	EMSRb
DF	Seq.	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B
0.8	3	EMSRb	EMSRb				
0.8	3	GVN	EMSRb	0.60%	-0.81%	0.33%	-0.30%
0.8	3	Netbid	EMSRb	-0.77%	-0.73%	-0.05%	0.22%
0.8	3	DAVN	EMSRb	1.14%	-0.55%	0.48%	-0.55%
0.8	3	HBP	EMSRb	0.12%	-1.03%	0.75%	-0.30%
0.8	3	GVN	GVN	0.03%	-0.12%	0.18%	0.29%
0.8	3	DAVN	GVN	0.21%	0.09%	0.01%	0.14%
0.8	3	DAVN	DAVN	0.26%	0.15%	0.14%	0.10%
0.8	3	DAVN	Netbid	0.75%	-1.01%	0.26%	-0.87%
0.8	3	HBP	GVN	0.00%	-0.05%	-0.26%	0.26%
0.8	3	Netbid	Netbid	-0.56%	-0.61%	-0.53%	-0.29%
0.8	3	DAVN	HBP	0.10%	-0.53%	0.17%	0.24%
0.8	3	Netbid	HBP	-0.96%	0.38%	-0.54%	0.53%
0.8	3	HBP	HBP	0.12%	0.14%	-0.21%	0.46%

Table 4.25: Network 1 vs. Network 2, DF = 0.8, Sequence 3

However, under Sequence 3, the results are very different. Table 4.25 shows the results of Network 1 vs. Network 2 under Demand Factor of 0.8, Sequence 3, with O-D control RM methods vs. EMSRb as well as competing head to head against each other. Under Sequence 3, passengers now have a choice between airlines. Local passengers flying between hubs and spokes could choose who they want to fly under Network 1, but not under Network 2. Connecting passengers flying from spoke to spoke could choose who they want to fly under both Networks. The effect of this difference can be seen from the percent change in revenue results under Sequence 3. Again, the percentage number presented in Table 4.25 represents the difference in revenue for each RM algorithm compared to the base case when both airlines use EMSRb. We can see that even at this low demand factor, most O-D control RM methods have a bigger effect on revenue in Network 1 than Network 2. This is because in Network 1, airlines compete for local traffic in addition to connecting traffic, making the decision of which passenger to accept more critical. This higher level of competition in Network 1, in most cases, amplifies the revenue effect of O-D control RM methods.

For head to head cases, a different picture emerged. Again, DAVN had a bigger positive revenue impact, and Netbid had a bigger negative revenue impact under Network 1. However, both GVN and HBP exhibited a smaller revenue impact under Network 1. This is because both GVN and HBP use unadjusted virtual classes based on total fares, such that the local fares end up in the lower virtual classes. If the seat a local passenger wants is not available, under Network 2 that passenger must either spill or sell-up. Under Network 1, however, that passenger could either spill or choose the other carrier, leading to a lower level of sell-up. As a result of a higher level of sell-up, GVN and HBP performed better under Network 2.

				Network 1		Network 2	
				% increase		% increase	
				compared to		compared to	
				EMSRb		EMSRb	
DF	Seq.	YM methods		EMSRb	EMSRb	EMSRb	EMSRb
		Airline A	Airline B	Airline A	Airline B	Airline A	Airline B
1.0	0	EMSRb	EMSRb				
1.0	0	GVN	EMSRb	0.88%		0.84%	
1.0	0	Netbid	EMSRb	0.63%		0.64%	
1.0	0	DAVN	EMSRb	0.98%		0.93%	
1.0	0	HBP	EMSRb	0.96%		0.90%	
1.0	2	EMSRb	EMSRb				
1.0	2	GVN	EMSRb	0.90%		0.89%	
1.0	2	Netbid	EMSRb	-0.20%		-0.28%	
1.0	2	DAVN	EMSRb	1.30%		1.31%	
1.0	2	HBP	EMSRb	1.17%		1.17%	

Table 4.26: Network 1 vs. Network 2, DF = 1.0, Sequences 0 and 2

Moving on to Demand Factor = 1.0, the average demand over all legs in the network equals the aircraft capacity of 100 seats. Again, because of the First-Choice-Only-Choice assumption under Sequences 0 and 2, we see no difference in RM algorithm performance between the two Networks, as expected, and as explained earlier.

				Network 1		Network 2	
				% increase		% increase	
				compared to		compared to	
				YM methods		EMSRb	EMSRb
DF	Seq.	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B
1.0	3	EMSRb	EMSRb				
1.0	3	GVN	EMSRb	0.37%	-0.30%	0.26%	-1.58%
1.0	3	Netbid	EMSRb	-0.64%	0.08%	-0.66%	-1.05%
1.0	3	DAVN	EMSRb	1.92%	-1.07%	1.46%	-2.03%
1.0	3	HBP	EMSRb	1.71%	-0.68%	1.32%	-1.78%
1.0	3	GVN	GVN	-0.11%	-0.10%	-0.67%	-0.80%
1.0	3	DAVN	GVN	1.09%	-1.25%	0.09%	-1.45%
1.0	3	DAVN	DAVN	0.44%	0.55%	0.00%	-0.21%
1.0	3	DAVN	Netbid	1.91%	-1.94%	1.15%	-1.85%
1.0	3	HBP	GVN	0.27%	-0.74%	-0.27%	-0.76%
1.0	3	Netbid	Netbid	-0.84%	-0.86%	-1.23%	-1.18%
1.0	3	DAVN	HBP	0.94%	-0.03%	-0.04%	-0.56%
1.0	3	Netbid	HBP	-1.32%	0.16%	-1.79%	-0.13%
1.0	3	HBP	HBP	0.19%	0.08%	-0.38%	-0.49%

Table 4.27: Network 1 vs. Network 2, DF = 1.0, Sequence 3

Table 4.27 presents the results from Sequence 3 under DF = 1.0. For the cases where O-D algorithms compete against EMSRb, again we see a bigger percent change in revenue under Network 1. This time, even HBP had a bigger revenue impact under Network 1. As for the cases where RM algorithms compete head to head with each other, we see DAVN exhibited a bigger revenue impact, while GVN, HBP, and Netbid exhibited a smaller revenue impact under Network 1. Again, this could be because with GVN and HBP, there was less sell-up under Network 1, leading to a smaller revenue improvement compared to the base case. As for Netbid, it exhibited negative revenue impacts under both networks, suggesting that it may still be affected by the small network size, and that it is letting on too many lower-fare paying passengers.

				Network 1		Network 2	
				% increase		% increase	
				compared to		compared to	
				EMSRb		EMSRb	
DF	Seq.	YM methods		EMSRb	EMSRb	EMSRb	EMSRb
		Airline A	Airline B	Airline A	Airline B	Airline A	Airline B
1.2	3	EMSRb	EMSRb				
1.2	3	GVN	EMSRb	-0.84%	0.65%	0.46%	-0.99%
1.2	3	Netbid	EMSRb	-0.94%	0.63%	-0.47%	-0.58%
1.2	3	DAVN	EMSRb	2.68%	-0.45%	2.44%	-1.67%
1.2	3	HBP	EMSRb	1.91%	-0.14%	2.47%	-1.56%
1.2	3	GVN	GVN	0.36%	0.56%	-0.15%	-0.16%
1.2	3	DAVN	GVN	2.69%	-0.32%	1.32%	-0.30%
1.2	3	DAVN	DAVN	1.39%	1.33%	0.79%	0.73%
1.2	3	DAVN	Netbid	2.90%	-1.03%	1.97%	-1.55%
1.2	3	HBP	GVN	1.70%	-0.25%	0.96%	-0.47%
1.2	3	Netbid	Netbid	-0.51%	-0.53%	-0.78%	-0.90%
1.2	3	DAVN	HBP	1.81%	1.03%	1.22%	0.91%
1.2	3	Netbid	HBP	-1.17%	0.63%	-1.48%	0.42%
1.2	3	HBP	HBP	1.30%	1.37%	0.60%	0.75%

Table 4.28: Network 1 vs. Network 2, DF = 1.2, Sequence 3

Table 4.28 shows the results for Sequence 3 at DF = 1.2. For O-D RM algorithms competing against EMSRb, again DAVN had a bigger revenue impact under Network 1. As for HBP, its result was inconsistent because it again shows a smaller revenue impact under Network 1. HBP showed a bigger revenue impact under DF=1.0, but showed a smaller revenue impact under DF = 0.8. As for GVN, an important result occurs here. Under Network 1, at such a high demand factor, GVN performed poorly compared to EMSRb. This is due to the algorithm being too greedy at a high demand factor, and accepting too many connection bookings. These connecting bookings in turn displaced the seats that would otherwise be available to two local passengers on the legs in and out of the hub. Under Network 2 however, we do not see such an effect even at this high level of demand. It may still occur at an even higher level of demand, but at DF = 1.2 under Network 2, GVN still performed better than EMSRb. This could be that under Network 1 where the two airlines compete for both local and connecting passengers, the airline using GVN becomes too greedy and thus spills its high-yield local passengers to the competing airline using EMSRb. This can be confirmed by the 0.65% improvement in revenue of the EMSRb airline at DF = 1.2. This effect cannot happen under Network 2 because the local passengers cannot be spilled to the competing airline. As a result, GVN

does not suffer as badly from being too greedy when the airlines have their own captive local passengers.

Furthermore, at this high demand factor, we see a reversal of results when both airlines use O-D RM algorithms. DAVN still had a bigger revenue impact under Network 1, same as with a smaller demand factor. Netbid still performed poorly with negative revenue impacts under both Networks. GVN and HBP, however, had a bigger revenue impact under Network 1, which is contrary to our previous results under lower demand factors. At this higher demand factor, when both airlines used GVN under Network 1, they did not suffer from being too greedy. Instead of spilling the local passengers to the competitor, since both airlines were short of seats, they forced more passengers to sell-up, leading to a higher level of revenue improvement under Network 1. Under Network 2, without the competitor also forcing their passengers to sell-up, both airlines suffered from being too greedy and simply spilled out their local passengers. As for HBP, since it is very similar to GVN, the same feedback effect existed and thus caused a bigger revenue impact under Network 1.

4.3.4 Summary

In this section, comparisons were made between two networks: Network 1 has both airlines sharing the same hub and competing for both local and connecting traffic, and Network 2 has both airlines having their own hubs and competing only for connecting traffic.

Under Sequences 0 and 2 when we assumed passengers have their first choice of path/fare class as their only choice, since airlines operate as if they are independent airlines, the hub location has no effect on the simulated relative performance of RM algorithms. The revenue performance relative to the base case of EMSRb vs. EMSRb is statistically identical under both Networks. It is only under Sequence 3 where passengers can choose alternative paths and fly on the competitor that whether there is competition in the local markets makes a difference in O-D RM algorithm performance.

GVN had a bigger positive revenue impact under Network 1 at $DF=0.8$ and 1.0 . This is because under Network 1 both airlines compete not only for connecting but also for local traffic. This higher level of competition makes the decision of which passenger to take more critical, leading to a larger revenue improvement for GVN under Network 1. At a demand factor of 1.2 , when only one airline used GVN while the other one used EMSRb, GVN became too greedy under Network 1 and its revenue suffered. Under Network 2, because local passengers could not fly on the competitor, being greedy did not hurt GVN as much. In the cases when both airlines used GVN, the opposite result occurred. Under Network 1, because the competitor also had limited seats for spilled local passengers, this forced the passengers to sell-up instead of spilling to the competitor, which allowed both airlines to benefit from using GVN. Under Network 2 when both airlines were too greedy, the local passengers they spilled could not be forced to sell-up, and both airlines suffered from the loss of the spilled passengers.

Netbid with only four spoke cities still suffers from the small network effect. There is too large of a gap between bid-price cut-off points, and thus too many low fare passengers fly. This caused Netbid to have a negative revenue impact under Sequence 3 at all demand factors for both Networks.

DAVN performed very well under all cases. It has positive revenue impacts under both networks, but caused a bigger improvement in revenue under Network 1. Similar to GVN, this is because under Network 1, airlines compete for both local and connection traffic. This higher level of competition makes the decision of which passenger to take more critical. The lack of competition on local paths under Network 2 resulted in the smaller revenue impact from DAVN.

Finally for HBP, at low demand factors, it had a bigger positive revenue impact under Network 1. Again, this was because of the higher level of competition for both local and connecting passengers under Network 1, which amplified the revenue effect of HBP. At

the higher demand factor, HBP did not suffer from being too greedy like GVN, and performed very well compared to other O-D control Revenue Management methods.

4.4 Different Implementations of DAVN and HBP O-D Seat Inventory Control Methods

4.4.1 Motivation

In the above experiments, it is clear that HBP and DAVN consistently performed better than the other O-D control RM methods. DAVN is an interesting algorithm to study because it is very robust, performing well under all conditions tested. HBP is another attractive algorithm because it performs almost as well as DAVN under most cases, without having to move to an O-D path based data collection and forecasting system. Since existing data and forecasts could be used under HBP, it is a much easier upgrade path for the airlines while still getting very good performance. It is important to study DAVN and HBP further because of their robust performance, and this section will focus on different implementations of HBP and DAVN to investigate how to make HBP and DAVN perform even better.

In Section 4.4.2, different d-factors will be explored and compared under HBP to see the sensitivity of changing d-factors is to the performance of the RM algorithm. In Section 4.4.3, different implementations of DAVN will be explored. These include DAVN without displacement cost to see the effect of displacement cost, different virtual class boundary schemes, and finally different pseudo fare and virtual class boundary revision schemes.

4.4.2 Heuristic Bidprice (HBP) Implementations

Under Heuristic Bidprice, a connecting (2-leg) booking is accepted when the following two conditions are met:

$$\text{FARE} > \text{EMSR}_{\text{leg1}} + \text{d-factor} * \text{EMSR}_{\text{leg2}}$$

and

$$\text{FARE} > \text{EMSR}_{\text{leg2}} + \text{d-factor} * \text{EMSR}_{\text{leg1}}$$

The d-factor is a factor related to the proportion of local/connecting passenger mix, and $\text{EMSR}_{\text{leg1}}$ and $\text{EMSR}_{\text{leg2}}$ are the EMSR values of the last seat on the respective legs. In this section, three different values of d-factors – 0.15, 0.40, and our base case 0.25 are compared under Network 3 to see the effect changing d-factors has on the revenue performance of HBP. The results are presented in Table 4.29.

DF	d-factor	Seq.	YM methods		% increase compared to	
					EMSRb	EMSRb
			Airline A	Airline B	Airline A	Airline B
1.0	0.15	3	HBP	EMSRb	1.11%	-1.85%
1.0	0.15	3	HBP	GVN	-0.71%	-1.05%
1.0	0.15	3	HBP	DAVN	-0.56%	-0.23%
1.0	0.15	3	HBP	Netbid	0.43%	-1.17%
1.0	0.15	3	HBP	HBP	-0.49%	-0.73%
1.0	0.40	3	HBP	EMSRb	1.83%	-2.24%
1.0	0.40	3	HBP	GVN	0.19%	-0.78%
1.0	0.40	3	HBP	DAVN	0.03%	-0.16%
1.0	0.40	3	HBP	Netbid	0.84%	-1.28%
1.0	0.40	3	HBP	HBP	0.00%	-0.15%
1.0	0.25	3	HBP	EMSRb	1.77%	-1.69%
1.0	0.25	3	HBP	GVN	-0.77%	-0.79%
1.0	0.25	3	HBP	DAVN	-0.59%	0.24%
1.0	0.25	3	HBP	Netbid	-0.96%	-0.82%
1.0	0.25	3	HBP	HBP	-0.47%	-0.33%

Table 4.29: HBP with d-factors at 0.15, 0.40, and 0.25

A d-factor = 0.25 was the base case that was used for all other HBP experiments. From the above table, we can see that by increasing the d-factor, the performance of HBP

increased as well. When Airline A used HBP while Airline B used EMSRb, the revenue performance of Airline A compared to base case improved from 1.11% to 1.77%, and finally to 1.83%, as we increase the d-factor from 0.15 to 0.25, and finally to 0.40. As for HBP vs. other RM algorithms, HBP performed about equally well under d-factors of 0.15 and 0.25 except for when it was against Netbid. However, under 0.40 it performed significantly well against Netbid. Airline A using HBP with d-factor of 0.40 achieved a revenue improvement over all other algorithms, including DAVN. This shows that HBP can also perform very well under the right settings.

To conclude, the airline using HBP performed well under d-factors of 0.15 and 0.25 when competing against leg based EMSRb algorithm, but it performed very well under a d-factor of 0.40. Therefore, d-factor values have a large impact on the performance of HBP, and it is important to investigate the most appropriate d-factors for different network scenarios.

4.4.3 DAVN Implementations

As DAVN was the best performing O-D control RM method under all conditions tested thus far, it would be interesting to investigate alternative implementations to further improve the robustness and performance of DAVN. This section will report on investigations into several experiments, including DAVN without displacement cost, different virtual class boundary methods, and finally different re-optimization schemes. In this section, first the DAVN without displacement cost experiment will be discussed. This is done to further understand the relative impacts of network optimization and ODF path-based forecasting on the DAVN algorithm. Next, virtual class boundary method and re-optimization scheme experiments will be presented together as further investigation into implementations of making DAVN perform better.

DAVN Without Displacement Cost

The DAVN without displacement cost experiment was run in order to investigate how much of the benefit of DAVN came from forecasting on a path/fare class level, and how

much of the benefit of DAVN came from the combination of path/fare class level forecasting as well as network displacement cost. Recall from Section 2.3.3 the workings of DAVN. Data collection and forecasting are done for each path/fare class. Virtual class forecasts for each leg are then constructed from path/fare class history by adding means and variances. Finally, for optimization and availability control, shadow prices for each leg are calculated in order to generate pseudo fares for each leg. These pseudo fares are then used in the EMSRb calculation for seat inventory control. The DAVN without displacement cost implementation, or DAVN(14), incorporates the path/fare class data collection and forecasting routines of DAVN, but does not perform the network displacement cost and pseudo fare calculations. DAVN(14) is essentially GVN with path based forecasting, as opposed to leg/virtual bucket forecasting. By comparing DAVN(14) results with the full DAVN results, we can find out how much of the improvement of DAVN came from the path forecasting component, and how much of the improvement came from the displacement cost, pseudo fare availability control.²

Table 4.30 shows the results of DAVN vs. DAVN(14) experiment under Network 1. Again, DAVN(14) results are with path/fare class forecasting only, while DAVN results added on the effect of displacement cost. The comparison between DAVN(14) and EMSRb shows us the portion of revenue gain from path forecasting, and the difference between DAVN and DAVN(14) shows us the portion of revenue gain from using the network displacement costs and pseudo fares. At demand factor = 0.8, DAVN and DAVN(14) performed equally well under Sequences 0 and 2. This indicates that all the revenue gain at this low demand level and no passenger choice came mainly from path-based forecasting. Even with passenger choice under Sequence 3, DAVN(14) had a gain of 0.95% over EMSRb, while DAVN had a 1.14 gain over EMSRb. This tells us that at this low demand factor, the majority of the gain of DAVN was because of better forecasting, while only 0.19% of the gain was because of network displacement costs and pseudo fares.

² Special thanks to Matthew Berge of The Boeing Company for suggesting this experiment.

Moving onto a demand factor of 1.0, under Sequences 0 and 2 with no passenger choice, about 1% of the total revenue improvement was because of better forecasting. Only 0.13 and 0.24% came from the displacement costs and pseudo fare availability control. Under Sequence 3, DAVN(14) saw a revenue improvement of 0.90%, while DAVN saw a revenue improvement of 1.92%, or 1.02% more. This tells us that at this demand factor, about half of the benefits of DAVN came from the better path forecasting, while the other half came from the network displacement cost and pseudo fare calculations.

DF	Seq.	YM methods		% increase compared to		% difference between DAVN and DAVN(14)
				EMSRb	EMSRb	
		Airline A	Airline B	Airline A	Airline B	
0.8	0	DAVN	EMSRb	0.15%		
	2	DAVN	EMSRb	0.56%		
	3	DAVN	EMSRb	1.14%	-0.55%	
0.8	0	DAVN (14)	EMSRb	0.15%		0.00%
	2	DAVN (14)	EMSRb	0.56%		0.00%
	3	DAVN (14)	EMSRb	0.95%	-0.56%	-0.19%
1.0	0	DAVN	EMSRb	1.07%		
	2	DAVN	EMSRb	1.30%		
	3	DAVN	EMSRb	1.92%	-1.07%	
1.0	0	DAVN (14)	EMSRb	0.94%		-0.13%
	2	DAVN (14)	EMSRb	1.06%		-0.24%
	3	DAVN (14)	EMSRb	0.90%	-0.46%	-1.02%
1.2	0	DAVN	EMSRb	1.71%		
	2	DAVN	EMSRb	1.74%		
	3	DAVN	EMSRb	2.68%	-0.45%	
1.2	0	DAVN (14)	EMSRb	0.98%		-0.72%
	2	DAVN (14)	EMSRb	1.00%		-0.74%
	3	DAVN (14)	EMSRb	0.74%	0.56%	-1.94%

Table 4.30: DAVN(14) vs. DAVN under Network 1

Finally moving onto DF = 1.2, we can see that under Sequences 0 and 2 with no passenger choice, again about 1% of the benefit came from better forecasting, and 0.72 and 0.74% came from displacement cost and pseudo fares. However under Sequence 3, about 0.74% of the improvement came from better path forecasting, while 1.94% of the improvement came from displacement cost and pseudo fares.

It is apparent that ODF path/fare class forecasting consistently brings in a revenue improvement over EMSRb of slightly less than 1%, decreasing as demand factors increase. Displacement cost and pseudo fares on the other hand, brings in anywhere from 0.19% at DF = 0.8 to 1.94% at DF = 1.2. This is understandable because as the demand increases, there are fewer seats available, and the importance of displacement cost becomes more significant. We can conclude that DAVN performed well both because of the better forecasting as well as a better seat availability control method. Under all demand factors, better forecasting contributed 0.74 to 1% of revenue improvement, and gave DAVN its robustness. Furthermore, as demand factors rise, displacement cost and pseudo fare availability control become more significant, further improving the performance of DAVN.

Virtual Class Bucket Methods (VCBM) and Re-optimization Schemes (REOPT)

Now that we understand why DAVN performed so well compared to other O-D RM algorithms, we can investigate different implementations to further improve the performance of DAVN. The two different implementation schemes to be tested here include Virtual Class Bucket Methods (VCBM) and Re-optimization Schemes (REOPT). For this experiment, DAVN was run under both PODS6 and PODS7 passenger choice assumptions. Network 3 at a demand factor of 1.0 was used for all comparisons. Before we begin, let's look at a few definitions that will better explain the experiments:

	Virtual Class Bucketing Method (VCBM)	
Re-optimization scheme (REOPT)	Manual / Network wide	Automated Leg Specific
Once at start of booking process	DAVN(11)	DAVN(21)
Every Time Frame	DAVN(12)	DAVN(22)

Table 4.31: DAVN Implementation Nomenclature

The above labeling of DAVN consists of DAVN(VCBM, REOPT). We have two different VCBM implementations – a setting of 1 denotes our base case, which uses user-input virtual class boundaries for the entire network. A setting of 2 denotes that virtual class boundaries are determined automatically by demand equalization for each leg separately. As for REOPT, again we have two different implementations – a setting of 1 again is our base case, and it denotes that network displacement costs (shadow prices) and virtual class boundaries are established only once at the beginning of the booking process. A setting of 2 denotes that network displacement costs and virtual class boundaries (when applicable – only VCBM=2 requires it) are calculated at every single time frame, or booking period “checkpoint”. In these experiments, each time frame consists of 7 days at first, shortening to 3 days within 14 days of departure, for a total of 16 time frames, or 16 revision opportunities.

Under VCBM = 1, we have another set of experiments run with the virtual class boundaries being 0.8 times the original value. The rationale behind this scheme was that since the pseudo fares are lower than the actual published fares, we should manually adjust the boundaries downward to better match the pseudo fare distribution. The following labeling scheme is used:

VC	1	2	3	4	5	6
DAVN(1,REOPT)	1300	700	420	310	215	140
DAVN(.8)(1,REOPT)	1300	560	336	248	172	112

Table 4.32: Virtual Class Upper Boundaries

With the definitions explained, we are ready to discuss the experiment results. The experiments are broken down into two sections – VCBM results and REOPT results. First let’s look at VCBM results under both PODS6 and PODS7.

DF	Seq	YM methods		PODS Version				PODS Version			
				6I2		7B1		6I2		7B1	
				% increase compared to		% increase compared to		% diff between DAVN(21) & DAVN(11)		% diff between DAVN(21) & DAVN(11)	
				EMSRb	EMSRb	EMSRb	EMSRb	Airline A	Airline B	Airline A	Airline B
Airline A	Airline B	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B		
1.0	3	EMSRb	EMSRb								
		DAVN(11)	EMSRb	2.09%	-1.92%	2.33%	-1.20%				
		DAVN(11)	GVN	0.06%	-1.01%	0.95%	0.22%				
		DAVN(11)	DAVN(11)	-0.10%	-0.32%	0.80%	0.78%				
		DAVN(21)	EMSRb	2.21%	-2.12%	2.51%	-1.36%	0.12%	-0.20%	0.18%	-0.16%
		DAVN(21)	GVN	0.43%	-1.07%	0.97%	0.13%	0.37%	-0.07%	0.02%	-0.09%
		DAVN(21)	DAVN(21)	0.02%	-0.11%	0.75%	0.72%	0.12%	0.21%	-0.05%	-0.06%

Table 4.33: Different Virtual Class Boundary Methods

Under DF = 1.0, the demand equalized virtual classes (VCBM = 2) performed better than our user-input virtual class boundaries (VCBM = 1) under most cases, although the difference was very small. By looking at the right side of the table where the difference between DAVN(21) and DAVN(11) is displayed, we can see that DAVN(21) outperformed DAVN(11) in all but one case. In that one case, the difference was too small to be conclusive. This tells us that at DF = 1.0, demand equalized virtual class boundaries (VCBM = 2) makes DAVN perform slightly better.

Next set of results are DAVN(.8)(12) results, where the virtual class upper boundaries are reduced by 20% to better match the pseudo fare distributions. Under this set of results, the pseudo fares are recalculated at every timeframe checkpoint.

DF	Seq	YM methods		PODS Version				PODS Version			
				6I2		7B1		6I2		7B1	
				% increase compared to		% increase compared to		% diff between DAVN(12) & DAVN(.8)(12)		% diff between DAVN(12) & DAVN(.8)(12)	
				EMSRb	EMSRb	EMSRb	EMSRb	Airline A	Airline B	Airline A	Airline B
Airline A	Airline B	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B		
1.0	3	EMSRb	EMSRb								
		DAVN(12)	EMSRb	1.77%	-2.64%	2.34%	-1.19%				
		DAVN(12)	GVN	0.08%	-1.24%	1.00%	0.22%				
		DAVN(12)	DAVN(12)	-0.17%	-0.20%	0.83%	0.81%				
		DAVN(.8)(12)	EMSRb	2.96%	-1.94%	2.72%	-1.32%	1.19%	0.70%	0.38%	-0.13%
		DAVN(.8)(12)	GVN	0.89%	-1.13%	1.20%	0.10%	0.81%	0.11%	0.20%	-0.12%
		DAVN(.8)(12)	DAVN(.8)(12)	0.79%	0.60%	0.96%	0.93%	0.96%	0.80%	0.13%	0.12%

Table 4.34: DAVN(.8)(12) vs. DAVN(12) results

By looking at the right side of the table where the difference between DAVN(12) and DAVN(.8)(12) is displayed, we can see that by reducing the virtual class boundaries by

20% to better match the pseudo fares, DAVN performed up to 1.19% better under PODS6 assumptions. Even under the PODS7 passenger choice model assumptions with lower levels of sell-up, lowering the virtual class boundaries also caused DAVN to perform up to 0.38% better when compared against EMSRb. This result was much bigger than the result from just changing the virtual class boundary methods, and thus confirmed that by better matching the virtual class boundaries and the pseudo fares, we can make DAVN perform significantly better.

Moving onto different Re-optimization schemes, Table 4.35 shows the comparison between Re-optimization at the beginning (REOPT = 1) and Re-optimization at every time frame (REOPT = 2) under both Virtual Class Boundary Methods.

DF	Seq	YM methods		PODS Version				PODS Version			
				6I2		7B1		6I2		7B1	
				% increase compared to		% increase compared to		% diff between REOPT=1 and REOPT=2		% diff between REOPT=1 and REOPT=2	
				EMSRb	EMSRb	EMSRb	EMSRb	Airline A	Airline B	Airline A	Airline B
Airline A	Airline B	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B	Airline A	Airline B		
1.0	3	EMSRb	EMSRb								
		DAVN(11)	EMSRb	2.09%	-1.92%	2.33%	-1.20%				
		DAVN(11)	GVN	0.06%	-1.01%	0.95%	0.22%				
		DAVN(11)	DAVN(11)	-0.10%	-0.32%	0.80%	0.78%				
		DAVN(12)	EMSRb	1.77%	-2.64%	2.34%	-1.19%	-0.32%	-0.72%	0.01%	0.01%
		DAVN(12)	GVN	0.08%	-1.24%	1.00%	0.22%	0.02%	-0.23%	0.05%	0.00%
		DAVN(12)	DAVN(12)	-0.17%	-0.20%	0.83%	0.81%	-0.07%	0.12%	0.03%	0.03%
		DAVN(21)	EMSRb	2.21%	-2.12%	2.51%	-1.36%				
		DAVN(21)	GVN	0.43%	-1.07%	0.97%	0.13%				
		DAVN(21)	DAVN(21)	0.02%	-0.11%	0.75%	0.72%				
		DAVN(22)	EMSRb	2.00%	-1.98%	2.45%	-1.21%	-0.21%	0.14%	-0.06%	0.15%
		DAVN(22)	GVN	0.17%	-1.08%	1.01%	0.27%	-0.26%	-0.01%	0.04%	0.14%
		DAVN(22)	DAVN(22)	0.41%	0.19%	0.94%	0.91%	0.39%	0.30%	0.19%	0.19%

Table 4.35: Different Re-optimization Schemes

When the existing user-input virtual class boundaries (VCBM = 1) were used, revisions at every time frame did not affect the revenue outcome very much. Under the PODS6 passenger choice assumption, DAVN(11) performed slightly better than DAVN(12). However, under the PODS7 passenger choice model with less sell-up, the two methods performed statistically the same.

On the other hand, when the demand equalized virtual class boundaries (VCBM = 2) were used, revisions at every time frame actually decreased the revenue performance of DAVN. Under the PODS6 assumption with less sell-up, DAVN without re-optimization (DAVN(21)) performed better than DAVN with re-optimization (DAVN(22)) except when both airlines used DAVN, in which case both airlines with re-optimization showed a large revenue gain. Similarly under PODS7, there was no statistical difference between DAVN(21) and DAVN(22), except when both airlines used DAVN. In that case, when both airlines used DAVN with re-optimization DAVN(22), they both performed 0.19% better than when both airlines used DAVN without re-optimization (DAVN(21)).

To conclude, for the DAVN experiments, first we experimented with DAVN(14) vs. DAVN. We learned that roughly 0.74 % to 1% of DAVN's revenue advantage came from the path/fare class forecasting approach used. This effect decreases slightly as the demand factor increased. On the other hand, the effect of using the displacement costs and pseudo fares for availability control became stronger as demand factors increase, understandably because more displacement takes place at higher demand factors. These two factors combined together gave DAVN its robust revenue advantage under all conditions.

In order to investigate implementations to improve DAVN's performance, different Virtual Class Boundary Methods (VCBM) and Re-optimization Schemes (REOPT) were tried. We found out that by using demand equalized virtual class boundaries, DAVN only performed slightly better than the base case, which used user-input virtual class boundaries. However, by manually reducing the user-input virtual class boundaries to better match the pseudo fares, we saw a very significant revenue improvement. This tells us that sophisticated Virtual Class Boundary Methods might not be as important as good matching of the boundaries to the pseudo fares.

Finally, different re-optimization schemes were experimented with. We found that with fixed input virtual class boundaries, revising the pseudo fares at every time frame does not affect the revenue outcome very much. However, with demand equalized virtual

class boundaries, revising both the boundaries and pseudo fares at every time frame may actually decrease the revenue outcome. The only exception was when both airlines revise their boundaries and pseudo fares at every time frame, and in that case both airlines gain in revenue. This result suggests that more frequent revisions of the pseudo fares and virtual class boundaries did not improve DAVN's revenue performance.

This concludes the Results and Discussions Chapter. In this chapter, questions about O-D based Revenue Management algorithms and implementations were answered through discussion of experiment results. In the next chapter, a conclusion of this thesis will be given to re-cap the discussions and findings.

Chapter 5

Conclusion

5.1 Summary of PODS Investigation Results

In this thesis, the results from the investigation of competitive impacts of Origin-Destination control were presented. The simulation tool used for this investigation is called the Passenger Origin Destination Simulator (PODS), and the details about this computer simulator can be found in Chapter 2. The network information as well as system settings for PODS used for the investigation are presented in Chapter 3, and finally the results of the investigation were discussed in Chapter 4.

In Chapter 1, five questions regarding Origin-Destination control were raised. The first question was whether network Origin-Destination control methods would show a revenue advantage over leg-based EMSRb fare class control. The second question was what were the impacts of passenger choice assumptions. The third question was the effect of competing airlines with co-located hubs vs. separate hubs. The final question was to look at the revenue differences of different implementations of some of the RM methods. In this thesis, experiments were run using PODS to find out the answer to the above four questions regarding network Origin-Destination control. The results were discussed in Chapter 4, and I'll summarize the findings in the rest of this section.

First, the revenue as well as load results from Origin-Destination control Revenue Management methods were compared against EMSRb fare class control. The four O-D methods compared here included Greedy Virtual Network (GVN), Network Bid-Price (Netbid), Displacement Adjusted Virtual Nesting (DAVN), and Heuristic Bid-Price (HBP). In terms of revenue, DAVN consistently performed the best when competing

against EMSRb. HBP performed as well as DAVN in some but not all cases, but it has the advantage that the existing leg-based forecasts can be used. GVN's performance was behind DAVN and HBP, but was still well above Netbid's performance. At high demand factors, GVN performed not as well because the algorithm became too "greedy", and took on too many connecting passengers who displaced high yield local passengers. Finally, Netbid performed the worst out of the four O-D control methods. In many cases, it performed even worse than base case EMSRb fare class control. The first reason for this was that the networks used for the investigations were small, resulting in big gaps between the deterministic bid-prices and thus too many low fare passengers ended up taking up seats from high fare passengers. This is the "small network effect". The second reason for Netbid's poor performance was that Netbid did not work well with pick-up forecasting and booking curve detruncation. With other combinations of forecasting and detruncation methods as well as a bigger network, Netbid may perform as well as the other O-D methods.

Furthermore, we also noticed that revenue gains by the airline using O-D control RM methods were accompanied by revenue losses on the part of the competing airline still using EMSRb fare class control. Some of the O-D revenue gains came at the expense of the second airline, exhibiting some "zero-sumness" in the simulated competitive environment.

As to why the O-D control RM methods performed better than EMSRb, we looked at leg and path passenger loads. We found that in terms of leg loads, the airline using O-D control RM methods carried more passengers on the long-haul, low-demand legs. It also carried a better passenger fare class mix on the high and medium demand legs. As for the competing EMSRb airline, its long-haul, high fare passengers were taken away by the airline using O-D control methods, leading to a loss in revenue. In terms of path loads, the airline using O-D control methods carried more connecting passengers on the bottleneck high demand legs. Most of these passengers travel on paths that traverse a high demand and a low demand leg. The O-D control airline carried these high-fare connecting passengers who could not find a seat under EMSRb fare class control. This

also led to a better (i.e. higher-yield) fare class passenger mix. Therefore by carrying a better fare class passenger mix as well as carrying more high-fare passengers, the airline using O-D control RM methods was able to have a better revenue performance.

The second question raised was what were the impacts of passenger choice assumptions. First, we tested different levels of demand correlation and whether passengers had a choice or not in the simulation. Then, we also tested a different passenger choice model, with a different willingness-to-pay formulation. This new formulation results in a lower level of sell-up. The findings of the impacts of passenger choice assumptions are summarized here.

As for different levels of demand correlation and passenger choice, we noticed different effects for different O-D control RM methods. GVN results were inconsistent as we increased the level of demand correlation and introduced passenger choice. Under some demand factors GVN results improved, while under others it decreased. As for Netbid, its revenue results were always positive with no demand correlation and no passenger choice, but once these assumptions were introduced, Netbid revenue results suffered. Finally, for DAVN and HBP, their performance relative to EMSRb fare class control improved as demand correlation and passenger choice were introduced into the assumptions.

Moving onto the new PODS passenger choice model with a different willingness-to-pay formulation and a lower level of sell-up, we noticed that the relative RM rankings stayed the same. We also observed that under the new model, the revenue impacts of O-D control RM methods versus EMSRb fare class control was also greater because EMSRb was most affected by the lower level of sell-up. Y and M class passenger loads were lower, but Q and B class loads were higher, combining to result in a slightly lower passenger load with the new model. This lower passenger load was due to the lower level of sell-up, leading to more passengers being spilled. This also led to a lower level of absolute revenue when under the new passenger choice model.

The third question was the effect of competing airlines with co-located hubs vs. separate hubs. Here we noticed different results from different O-D control RM methods under different demand factors. First, GVN results suffered at high demand factors because it was too greedy. The phenomenon was more apparent when the two hubs are co-located, because the spilled local passengers simply traveled on the competitor and improved the competitor's revenue. This phenomenon was less obvious when the two airlines had separate hubs because they did not compete with each other on local passengers. Next, for Netbid, it was still suffering from the "small network effect" under both networks, thus leading to poor revenue performance. As for DAVN, it performed very well under both networks. It had a better performance under the co-located hub scheme because airlines compete for both local and connecting passengers, and with DAVN's good performance, the ability to trade off more local and connecting passengers led to better revenue results. Finally for HBP, since HBP is very similar to GVN, its performance was also very similar to GVN at low demand factors under the two networks. However, unlike GVN, it did not suffer from being too greedy at higher demand factors, thus performing very well compared to other O-D control RM methods.

Finally the fourth question was to look at the revenue differences of different implementations of selected RM methods. In particular, for HBP, different d-factors were experimented. For DAVN, different Virtual Class Boundary Methods and different Re-optimization Schemes were investigated. For HBP, we noticed that the airline using HBP performed better with d-factors of 0.40, but it still performed very well with lower d-factors. D-factor values have a large impact on the revenue performance, so it is important to investigate the most appropriate d-factors for different network scenarios. As for DAVN, under different Virtual Class Boundary Methods, we noticed that demand equalized virtual class boundaries had only a small positive revenue benefit when compared to user-input virtual class boundaries. When the virtual class boundaries were lowered manually to better match the pseudo fares, DAVN's performance improved, suggesting that demand equalized virtual class boundaries are not as important as better matching of the boundaries to the pseudo fares. As for different re-optimization schemes under DAVN, we found that by re-optimizing virtual class boundaries and pseudo fares at

every time frame, it may actually decrease the revenue performance. There was no significant benefit under DAVN to revise virtual class boundaries and pseudo fares frequently.

This concludes the summaries of investigations presented in this thesis. In the next section, suggestions for future research directions will be presented.

5.2 Future Research Directions

From the findings in this thesis, we can see several areas for future research:

For Netbid, it suffers from the “small network effect” as well as incompatible forecasting and detrunctation methods. One area for future research would be to move onto an even bigger network to see if Netbid performs better. This is actually a step being taken for PODS research in the next year.

We have also seen exciting results from the DAVN experiments. Future research can be done to investigate why re-optimization does not improve revenue results. We can also investigate other alternative virtual class bucketing methods to improve the performance of DAVN even further.

As for HBP experiment results, we noticed that d-factors have a large impact on HBP revenue results. Another possible area of future investigation is to experiment and try to understand d-factors better. Also, we can experiment with different networks to see if we can explain how d-factors interact with network topology.

Finally, we can further tap into PODS’ capability by introducing overbooking and cancellations. Up to this point, all experiments were run with no overbooking and no cancellation, and it would be interesting to see how overbooking and cancellation would affect the performance of O-D control RM methods.

The above are just some ideas of what we can do to further understand the competitive impacts of Origin-Destination control. This is by no means a complete list. With airlines adopting O-D control RM methods, more research is needed to further understand the implications and effects on network revenue. This thesis is just the tip of an iceberg, and hopefully we can continue doing research on this area.

Bibliography

- Belobaba, P. P. 1987. Air Travel Demand and Seat Inventory management. MIT Flight Transportation Laboratory Report R87-7
- Skwarek, D. K. 1996. Competitive Impacts of Yield Management System Components: Forecasting and Sell-up Models. MIT Flight Transportation Laboratory Report R96-6
- Wei, Y. 1997. Airline O-D Control Using Network Displacement Concepts.
- Williamson, E. L. 1992. Airline Network Seat Inventory Control: Methodologies and Revenue Impacts. MIT Flight Transportation Laboratory Report R92-3
- Wilson, J. L. 1995. The Value of Revenue Management Innovation in a Competitive Airline Industry. MIT Flight Transportation Laboratory Report R95-8
- Zickus, J. S. 1998. Forecasting for Airline Network Revenue Management: Revenue and Competitive Impacts

Appendix A

PODS Input File Descriptions

© The Boeing Company, 1997

PODS6I INPUT

READ(1,1000) DESC

description of case (1 line)

READ(1,*,END=100) NAL,NLEG,NMKT,NPTYP,NFCLS,NRST
,NBCRV,NTF,NOB,NBURN,NSAMP,NTRIAL

nal = # airline
 nleg = # legs
 nmkt = # markets
 nptyp = # pax types
 nfcls = # fare classes
 nrst = # restriction categories
 nbcrv = # booking curves
 ntf = # time frames
 nob = # observations used in YM forecaster
 nburn = # samples burned before statistics collected
 nsamp = total # samples (days)
 ntrial = # trials

READ(1,*) SKF,MKF,TKF,CKF

skf = system k-factor
 mkf = market k-factor
 tkf = pax type k-factor
 ckf = attributed cost k-factor

READ(1,*) NEWSD

newsd = flag for new (=1) random number seed

READ(1,*) IGAMMA

igamma = whether (=1) gamma or normal (=0) demand generator

READ(1,*) (IDBD(N),N=1,NTF)

idbd(n) = days before departure, start of time frame n

READ(1,*) FCOC

fcoc = flag whether (≥ 1) or not ($=0$) first pax choice is only choice
 [note: fcoc = 2 implies perfect history]

DO T=1,NPTYP

READ(1,*) ZF1(T),ZF2(T)

READ(1,*) GOINT(T),GOSLP(T)

READ(1,*) UFAINT(T),UFASLP(T)

READ(1,*) CPQIINT(T),CPQISLP(T)

READ(1,*) CRPNINT(T),CRPNSLP(T)

DO R=1,NRST

READ(1,*) CRSTINT(T,R),CRSTSLP(T,R)

ENDDO

ENDDO

zf1(t) = primary z-factor, pax type t

zf2(t) = secondary z-factor, pax type t

goint(t), goslp(t) = intercept, slope of max-willing-to-pay, pax type t

ufaint(t), ufasp(t) = intercept, slope of cost assigned

to unfavorite airlines, pax type t

cpqint(t), cpqislp(t) = intercept, slope of cost assigned to unit of

path quality index, pax type t

crpnint(t), crpnslp(t) = intercept, slope of cost assigned

to replanning, pax type t

crstint(t,r), crstslp(t,r) = intercept, slope of cost assigned

to restriction category r, pax type t

DO B=1,NBCRV

READ(1,*) (PBOOK(B,TF),TF=1,NTF)

ENDDO

pbook(b,tf) = (cum) probability of booking, booking curve b,

time frame tf [NOTE: pbook(b,ntf) = 1.]

READ(1,*) CNCPEN

DO B=1,NBCRV

READ(1,*) (PCANC(B,J),J=0,1)

ENDDO

cncpen = amount of cancellation penalty, given an airline designates
 a fare class to employ such a penalty

pcanc(b,j) = daily probability of cancellation, booking curve b, given

j = 0: no cancellation penalty, j = 1: cancellation penalty

DO B=1,NBCRV

READ(1,*) (PSHOW(B,J),J=0,1)

ENDDO

pshow(b,j) = probability of show, booking curve b, given

j = 0: no no-show penalty, j = 1: no-show penalty

```

DO A=1,NAL
  DO F=1,NFCLS
    READ(1,*) LTFFC(A,F),IPENC(A,F),IPENNS(A,F)
      ,(IDRST(A,F,R),R=1,NRST)
  ENDDO
ENDDO

```

ltffc(a,f) = last time frame fare class f is available, airline a
 ipenc(a,f) = flag indicating whether (=1) or not (=0) cancellation
 penalty applies to fare class f for airline a
 ipenns(a,f) = flag indicating whether (=1) or not (=0) fare is retained
 by airline in event of pax no-show, fare class f, airline a
 idrst(a,f,r) = flag indicating whether (=1) or not (=0)
 restriction r applies to fare class f, airline a

```

DO A=1,NAL
  READ(1,*) METHYM(A),METHNST(A),METHVRT(A),BPSCL(A),HBPCON(A)
    ,ISUP(A),METHFC(A),METHTRC(A),METHFR(A),ZOB(A)
ENDDO

```

methym(a) = index of YM optimization method, airline a
 leg based: 1 = fcfs, 2 = emsr-a, 3 = emsr-b
 dynamic vn: 12 = emsr-a, 13 = emsr-b, 14 = emsr-b (no displ)
 heuristic bp: 22 = emsr-a, 23 = emsr-b
 formal bp: 31 = network optimization
 methnst(a) = whether 'theft' (=1) or standard (=2) nesting used
 methvrt(a) = whether (=1) or not (=0) airline a uses virtual classes
must = 1 for methym = 10 - 29
must = 0 for methym = 31
 bpscl(a) = bid-price scaling, airline a
used only if methym = 10-19, 31
 hbpcn(a) = bid-price calc constant, airline a
used only if methym = 22, 23
 isup(a) = flag whether (=1) or not (=0) non-zero prob sell-up, airline a
used only with emsr-b (methym = 3, 13, 23)
 methfc(a) = index of forecasting method, airline a
 1 = BIH vs BAD regression
 2 = pick-up, 3 = efficient
 methtrc(a) = index of detruncator used, airline a
 0 = no detrunc, 1 = prob, 2 = proj,
 3 = simple pickup, 4 = local BC pickup
 methfr(a) = index of leg fare calculation method, airline a
 1 = non-stop path, 2 = dist wt avg of all paths
 3 = wt avg of reduced (by dsplcmnt) path fare
must = 3 for methym = 2 & 3 if methvrt = 1
3 implied for methym = 10 - 29
 zob(a) = sigma scaling for booking capacity calc, airline a

```
DO A=1,NAL
  IF(METHYM(A).GE.20) THEN
    READ(1,*) IBLV(A),NREOP(A)
  ENDIF
ENDDO
      iblv(a) = flag indicating for heuristic and network bidprice availability
processor whether (=1) or not (=0) forecast believed, airline a
      implication: forecast decremented by bookings if iblv = 1
      nreop(a) = # bookings between executions of avail processor, airline a
```

```
DO A=1,NAL
  IF(METHYM(A).GE.10.AND.METHYM(A).LT.20) THEN
    READ(1,*) METHDSP(A),METHRVC(A)
  ENDIF
ENDDO
      methodsp(a) = displacement method, airline a
      0 = no displ, 1 = input boundaries, 2 = dem equalized boundaries
      methrvc(a) = virtual class revision scheme, airline a
      1 = start of time frame one, 2 = every time frame
```

```
DO A=1,NAL
  IF(METHVRT(A).EQ.1) THEN
    READ(1,*) NVCAL(A)
    DO V=1,NVCAL(A)
      READ(1,*) FAREX(A,V)
    ENDDO
  ENDIF
ENDDO
      nvcal(a) = # virtual classes, airline a
      farex(a,v) = upper limit, virtual class v, airline a
```

```
DO A=1,NAL
  IF(ISUP(A).EQ.1) THEN
    DO V=1,NVCAL(A)
      READ(1,*) PSUP(A,V)
    ENDDO
  ENDIF
ENDDO
      psup(a,v) = assumed prob of sell-up, virtual class v, airline a
```

```
DO L=1,NLEG
  READ(1,*) LAL(L),CAP(L),DIST(L),LCAT(L)
ENDDO
      lal(l) = airline operating leg l
      cap(l) = capacity airplane used on leg l
      dist(l) = distance (SM) leg l
      lcat(l) = estimation category associated with leg l
```

```
DO M=1,NMKT
  READ(1,*) INSTC,INTOD,NPM,IRBIN,HSTART,DELT,BFARE(M),DBPEN(M)
ENDDO
```

instc = flag indicating whether (=1) or not (=0)
 mean schedule tolerance is input
 [NOTE: if 0, standard curve used]
 intod = flag indicating whether (=1) or not (=0)
 allocated time-of-day curve is input
 [NOTE: if 0, standard curve used]
 npm = # paths in market
 irbin = range bin of market (1 = 0-500, 2 = 500 - 1000 ...)
 hstart = start hour of day (nominal = 4)
 delt = market delta-t (nominal = min dt-at in decimal hours)
 bfare(m) = base fare, market m
 dbpen(m) = denied boarding penalty, market m

```
IF(INSTC.EQ.1) THEN
  READ(3,*) (STCMU(T),T=1,NPTYP)
ELSE
  READ(3,*) (BFRAC(T),T=1,NPTYP)
ENDIF
IF(INTOD.EQ.1) THEN
  DO T=1,NPTYP
    READ(3,*) (DIN(I,T),I=1,24)
  ENDDO
ENDIF
```

stcmu(t) = mean schedule tolerance, pax type t
 bfrac(t) = business fraction, pax type t
 [NOTE: used only with standard curves]
 din(i,t) = input allocated demand, time i-.5, pax type t

```
READ(1,*) (DMU(M,T),T=1,NPTYP)
READ(1,*) (BCUSE(M,T),T=1,NPTYP)
DO A=1,NAL
  READ(1,*) PFAL(M,A),(FARE(M,F,A),F=1,NFCLS)
ENDDO
```

dmu(m,t) = mean demand, market m, pax type t
 bcuse(m,t) = booking curve used, market m, pax type t
 pfal(m,a) = prob airline a is a favorite in market m
 fare(m,f,a) = fare, market m, fare class f, airline a

```
DO P1=1,NPM
  P=P1+NPATH
  READ(1,*) IDT,IAT,IDALP(P),IPQI(P),NLP,(LEGIN(N),N=1,NLP),PCAT(P)
ENDDO
```

idt = path departure time (std form -- 1000, 1425...)
 iat = path arrival time (std form -- 1200, 1625...)
 idalp(p) = identity of airline operating path p
 ipqi(p) = path quality index, path p
 (nominally: 1 + #stops + #connects * 2)
 nlp = # legs used by path
 legin(n) = identity of nth leg used
 pcat(p) = estimation category associated with path p

Appendix B

A Sample PODS Input File

A sample input file for Network 1, with only one market listed

Appendix C

Passenger Loads and Class Distribution Under PODS6 vs. PODS7

PODS 7, scen 6b, EMSRb vs. EMSRb

LEG	AL	TOT	LOADS BY FARE CLASS			
			Y	M	B	Q
1	1	91.19	40.38	18.91	7.71	24.19
2	1	91.32	40.96	19.01	7.69	23.66
3	1	90.59	36.31	17.12	5.40	31.76
4	1	90.61	36.04	17.08	5.35	32.14
5	1	86.09	32.17	15.43	3.94	34.54
6	1	85.89	32.20	15.30	4.04	34.35
7	1	73.92	26.99	13.03	3.10	30.80
8	1	73.77	26.90	13.00	3.08	30.79
9	1	62.07	22.60	10.97	2.62	25.88
10	1	61.92	22.55	11.02	2.63	25.71
11	1	49.82	17.94	8.96	2.13	20.79
12	1	50.17	17.91	9.00	2.11	21.15
13	2	91.20	40.70	18.94	7.66	23.90
14	2	91.31	40.29	18.84	7.72	24.46
15	2	90.73	36.27	17.07	5.40	31.99
16	2	90.65	36.43	17.16	5.36	31.69
17	2	85.85	31.98	15.39	3.98	34.51
18	2	85.90	32.13	15.34	4.05	34.38
19	2	73.63	26.90	12.83	3.12	30.77
20	2	73.68	26.80	12.95	3.08	30.85
21	2	61.88	22.46	11.00	2.65	25.79
22	2	62.08	22.54	11.07	2.64	25.82
23	2	50.09	17.98	9.01	2.17	20.93
24	2	50.30	18.05	9.01	2.08	21.16

PODS6, scen 6b, EMSRb vs. EMSRb

LEG	AL	TOT	LOADS BY FARE CLASS			
			Y	B	M	Q
1	1	91.75	45.20	15.98	23.28	7.29
2	1	91.85	45.21	15.98	23.45	7.21
3	1	91.49	40.73	14.59	15.61	20.56
4	1	91.39	40.71	14.49	15.98	20.21
5	1	87.64	36.01	13.02	9.70	28.92
6	1	87.42	35.84	12.97	9.77	28.84
7	1	75.60	30.02	11.13	7.19	27.26
8	1	75.47	30.03	11.01	7.07	27.37
9	1	63.82	25.11	9.53	6.22	22.96
10	1	63.59	25.20	9.39	6.08	22.92
11	1	51.42	20.22	7.73	5.04	18.44
12	1	51.47	20.19	7.78	4.75	18.74
13	2	91.84	45.16	15.87	23.11	7.70
14	2	91.89	45.61	16.07	23.43	6.78
15	2	91.61	40.33	14.46	15.59	21.24
16	2	91.41	40.66	14.45	15.89	20.40
17	2	87.43	35.79	12.96	9.69	28.99
18	2	87.68	36.15	12.96	9.60	28.97
19	2	75.61	30.07	11.01	7.14	27.38
20	2	75.94	30.04	11.10	7.18	27.62
21	2	63.85	25.29	9.57	6.17	22.83
22	2	63.74	25.14	9.42	6.07	23.11
23	2	51.30	20.30	7.78	4.98	18.24
24	2	51.93	20.23	7.82	4.84	19.05

PODS 7, scen 6b, VEMSRb vs. EMSRb

LEG	AL	TOT	LOADS BY FARE CLASS			
			Y	M	B	Q
1	1	91.27	40.49	18.22	3.55	29.02
2	1	91.28	40.66	18.20	3.42	29.00
3	1	88.68	36.32	17.32	3.30	31.74
4	1	88.66	36.48	17.12	3.46	31.60
5	1	88.33	32.12	15.40	4.19	36.62
6	1	88.23	32.14	15.36	4.19	36.55
7	1	79.18	26.82	13.14	2.82	36.39
8	1	79.20	26.66	13.18	2.80	36.56
9	1	68.58	22.58	11.43	2.31	32.26
10	1	68.83	22.74	11.37	2.32	32.40
11	1	57.01	18.12	9.35	1.86	27.68
12	1	56.73	17.95	9.37	1.82	27.59
13	2	90.82	40.55	19.02	6.11	25.14
14	2	90.91	40.45	18.88	6.07	25.51
15	2	90.40	36.38	17.23	4.67	32.12
16	2	90.46	36.24	17.33	4.71	32.19
17	2	84.91	32.67	15.30	3.29	33.65
18	2	84.53	32.51	15.15	3.28	33.59
19	2	72.41	27.06	12.88	2.53	29.94
20	2	72.23	26.94	12.82	2.54	29.93
21	2	60.54	22.45	10.94	2.08	25.07
22	2	60.31	22.56	10.93	2.03	24.79
23	2	48.27	17.80	8.83	1.54	20.10
24	2	48.78	17.96	8.98	1.58	20.26

PODS6, scen 6b, VEMSRb vs. EMSRb

LEG	AL	TOT	LOADS BY FARE CLASS			
			Y	B	M	Q
1	1	92.70	45.00	13.47	8.47	25.76
2	1	92.69	45.14	13.38	8.52	25.65
3	1	89.50	40.43	13.81	6.83	28.44
4	1	89.63	40.42	13.77	7.05	28.40
5	1	89.33	35.60	12.84	10.96	29.93
6	1	89.45	35.58	12.78	10.98	30.11
7	1	81.07	29.59	11.04	5.80	34.64
8	1	80.80	29.26	10.98	5.74	34.82
9	1	70.73	24.87	9.65	4.74	31.47
10	1	71.12	25.00	9.67	4.71	31.74
11	1	60.66	20.13	7.98	3.52	29.03
12	1	59.99	20.05	8.00	3.41	28.53
13	2	90.65	45.47	16.43	17.65	11.09
14	2	90.59	45.61	16.50	17.68	10.80
15	2	90.78	40.52	14.71	12.00	23.55
16	2	90.90	40.59	14.86	12.16	23.29
17	2	86.07	36.20	13.09	6.96	29.81
18	2	85.71	36.07	12.98	6.89	29.77
19	2	73.67	30.14	11.05	4.88	27.60
20	2	73.49	30.06	10.96	4.82	27.66
21	2	61.24	25.06	9.39	3.93	22.86
22	2	61.03	25.12	9.34	3.83	22.74
23	2	48.23	20.03	7.60	2.49	18.11
24	2	48.94	20.04	7.65	2.49	18.76

PODS 7, scen 6b, Netbid vs. EMSRb

LEG	AL	TOT	LOADS BY FARE CLASS			
			Y	M	B	Q
1	1	93.87	38.81	16.32	3.41	35.32
2	1	93.88	39.38	16.34	3.38	34.78
3	1	91.99	34.74	15.28	3.23	38.73
4	1	91.90	34.56	15.33	3.21	38.80
5	1	89.74	31.22	14.69	2.79	41.05
6	1	89.72	30.91	14.62	2.88	41.30
7	1	79.50	26.11	12.90	2.48	38.02
8	1	79.67	26.19	12.94	2.45	38.09
9	1	69.39	22.30	11.25	2.20	33.63
10	1	69.21	22.13	11.23	2.18	33.67
11	1	56.90	17.85	9.31	1.92	27.83
12	1	56.63	17.90	9.32	1.92	27.50
13	2	90.34	41.27	19.81	7.01	22.25
14	2	90.56	41.39	19.74	6.94	22.50
15	2	90.07	37.06	17.89	5.42	29.70
16	2	90.17	36.90	17.99	5.45	29.82
17	2	84.93	32.64	15.73	3.72	32.84
18	2	85.16	32.63	15.80	3.76	32.96
19	2	71.98	27.09	13.01	2.63	29.25
20	2	72.11	27.30	13.11	2.60	29.10
21	2	60.26	22.76	11.03	2.09	24.38
22	2	60.02	22.58	10.90	2.04	24.49
23	2	48.31	18.02	8.88	1.61	19.80
24	2	48.41	18.05	8.97	1.56	19.84

PODS6, scen 6b, Netbid vs. EMSRb

LEG	AL	TOT	LOADS BY FARE CLASS			
			Y	B	M	Q
1	1	93.61	44.52	12.41	8.22	28.46
2	1	93.66	44.73	12.29	7.80	28.84
3	1	92.82	39.04	12.41	7.01	34.36
4	1	92.94	39.32	12.39	7.19	34.04
5	1	89.84	34.79	12.34	5.83	36.89
6	1	90.30	34.81	12.31	5.88	37.30
7	1	82.64	29.04	10.84	4.50	38.26
8	1	82.78	28.89	10.88	4.50	38.51
9	1	73.59	24.80	9.59	4.13	35.08
10	1	73.51	24.85	9.50	4.09	35.08
11	1	61.33	20.13	7.93	3.69	29.58
12	1	60.65	20.08	7.99	3.54	29.03
13	2	90.65	46.42	17.05	20.11	7.06
14	2	90.49	46.62	17.21	20.18	6.48
15	2	90.88	41.77	15.53	14.68	18.90
16	2	90.83	41.80	15.52	15.03	18.49
17	2	86.65	36.42	13.31	9.28	27.64
18	2	86.54	36.40	13.29	9.22	27.63
19	2	72.51	30.56	11.23	4.66	26.05
20	2	72.66	30.57	11.16	4.70	26.23
21	2	59.63	25.29	9.42	3.33	21.59
22	2	59.64	25.39	9.36	3.27	21.61
23	2	47.29	20.06	7.52	2.45	17.27
24	2	47.75	20.02	7.54	2.43	17.75

PODS 7, scen 6b, DAVN vs. EMSRb

LEG	AL	TOT	LOADS BY FARE CLASS			
			Y	M	B	Q
1	1	88.15	41.53	17.81	2.50	26.31
2	1	88.64	41.67	17.69	2.43	26.85
3	1	87.10	37.19	16.41	2.76	30.75
4	1	87.04	37.14	16.42	2.67	30.82
5	1	85.93	32.85	15.33	4.14	33.61
6	1	85.55	32.76	15.31	4.14	33.34
7	1	79.97	26.93	13.18	2.81	37.04
8	1	80.17	27.03	13.18	2.87	37.08
9	1	69.88	22.74	11.40	2.36	33.38
10	1	69.84	22.76	11.37	2.29	33.41
11	1	58.08	18.46	9.38	1.81	28.43
12	1	57.34	18.33	9.34	1.84	27.84
13	2	91.40	40.37	19.96	7.42	23.65
14	2	91.12	40.43	20.05	7.35	23.28
15	2	90.92	35.83	18.25	5.54	31.29
16	2	90.53	36.24	18.16	5.64	30.50
17	2	85.62	32.12	15.41	3.89	34.20
18	2	85.48	31.75	15.35	3.86	34.52
19	2	71.39	26.73	12.88	2.54	29.24
20	2	71.54	26.89	12.99	2.50	29.16
21	2	59.48	22.37	10.91	2.02	24.18
22	2	59.53	22.28	11.02	2.01	24.22
23	2	47.77	17.87	8.83	1.56	19.50
24	2	48.26	17.78	8.85	1.56	20.08

PODS6, scen 6b, DAVN vs. EMSRb

LEG	AL	TOT	LOADS BY FARE CLASS			
			Y	B	M	Q
1	1	88.78	46.82	12.98	4.92	24.06
2	1	89.08	47.18	12.67	5.05	24.18
3	1	86.81	41.92	12.88	5.28	26.74
4	1	86.70	41.59	12.76	5.32	27.02
5	1	85.87	36.85	12.76	9.74	26.53
6	1	85.96	37.07	12.88	9.87	26.13
7	1	82.80	30.09	11.26	6.15	35.30
8	1	83.11	30.19	11.17	6.09	35.66
9	1	73.83	25.42	9.82	4.61	33.97
10	1	74.27	25.50	9.78	4.58	34.41
11	1	62.26	20.51	7.98	3.43	30.33
12	1	61.83	20.43	8.02	3.31	30.07
13	2	91.42	45.19	17.71	21.96	6.56
14	2	91.43	45.35	17.90	21.65	6.53
15	2	91.68	40.64	16.15	16.08	18.81
16	2	91.55	40.41	16.14	15.83	19.17
17	2	87.78	35.69	13.28	10.60	28.20
18	2	87.99	35.84	13.21	11.05	27.89
19	2	72.35	30.05	10.93	4.96	26.40
20	2	72.28	29.91	11.14	4.94	26.29
21	2	59.37	25.06	9.28	3.52	21.50
22	2	59.37	25.16	9.27	3.48	21.46
23	2	47.04	19.92	7.64	2.30	17.19
24	2	47.42	19.89	7.48	2.38	17.68

PODS 7, scen 6b, HBP(NAP) vs. EMSRb

LEG	AL	TOT	LOADS BY FARE CLASS			
			Y	M	B	Q
1	1	91.78	40.64	18.70	3.44	29.00
2	1	92.04	40.44	18.64	3.43	29.52
3	1	90.12	36.40	17.04	3.18	33.49
4	1	90.10	36.52	16.95	3.17	33.46
5	1	87.40	32.43	15.18	4.16	35.64
6	1	87.38	32.44	15.24	4.19	35.51
7	1	77.79	26.93	13.03	2.99	34.84
8	1	77.60	26.76	13.06	3.03	34.76
9	1	69.73	22.47	11.31	2.24	33.72
10	1	69.81	22.45	11.20	2.23	33.93
11	1	57.23	18.04	9.32	1.81	28.06
12	1	57.03	18.09	9.37	1.77	27.79
13	2	90.63	40.51	19.13	6.52	24.47
14	2	91.14	40.41	19.17	6.33	25.23
15	2	90.19	36.29	17.34	4.94	31.63
16	2	90.28	36.29	17.30	4.94	31.75
17	2	85.36	32.21	15.36	3.55	34.23
18	2	84.81	32.08	15.22	3.57	33.94
19	2	72.65	26.95	12.91	2.65	30.14
20	2	72.18	26.92	12.95	2.64	29.67
21	2	60.10	22.41	10.90	1.99	24.79
22	2	59.72	22.45	10.86	1.99	24.43
23	2	48.31	17.96	8.85	1.55	19.95
24	2	48.84	17.99	8.96	1.58	20.32

PODS6, scen 6b, HBP(NAP) vs. EMSRb

LEG	AL	TOT	LOADS BY FARE CLASS			
			Y	B	M	Q
1	1	91.72	45.60	14.78	9.07	22.28
2	1	91.83	45.18	14.57	9.18	22.90
3	1	90.05	41.08	14.06	7.00	27.92
4	1	90.01	40.89	14.07	6.88	28.18
5	1	88.47	36.13	12.68	11.42	28.24
6	1	88.54	36.39	12.80	11.56	27.79
7	1	79.97	29.90	10.96	7.04	32.07
8	1	80.03	30.12	11.01	6.97	31.93
9	1	74.00	25.08	9.73	4.37	34.82
10	1	74.30	25.19	9.71	4.37	35.04
11	1	61.20	20.37	8.03	3.31	29.50
12	1	60.87	20.41	8.03	3.31	29.12
13	2	90.93	45.24	16.36	19.40	9.92
14	2	91.09	45.52	16.50	19.48	9.58
15	2	90.95	40.55	14.90	13.53	21.97
16	2	90.84	40.61	14.81	13.40	22.01
17	2	86.79	35.95	13.04	8.28	29.52
18	2	87.01	36.20	13.08	8.40	29.33
19	2	74.22	30.07	11.11	5.56	27.48
20	2	74.54	30.13	11.20	5.60	27.61
21	2	60.17	25.26	9.35	3.23	22.32
22	2	60.03	25.15	9.25	3.28	22.35
23	2	47.93	20.12	7.56	2.44	17.81
24	2	48.56	20.19	7.65	2.41	18.30