A MARKOV DECISION PROCESS MODEL FOR AIRLINE MEAL PROVISIONING

by

Jason Hidekazu Goto

B.A.Sc. Mechanical Engineering, University of Waterloo, 1995

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF MASTER OF SCIENCE (BUSINESS ADMINISTRATION)

in

THE FACULTY OF GRADUATE STUDIES

(Department of Commerce and Business Administration)

We accept this thesis as conforming To the required standard

THE UNIVERSITY OF BRITISH COLUMBIA

April 1999

© Jason Hidekazu Goto, 1999

In presenting this thesis in partial fulfilment of the requirements for an advanced degree at the University of British Columbia, I agree that the Library shall make it freely available for reference and study. I further agree that permission for extensive copying of this thesis for scholarly purposes may be granted by the head of my department or by his or her representatives. It is understood that copying or publication of this thesis for financial gain shall not be allowed without my written permission.

Department of COMMERCE - BUSINESS ADMINISPRATION

The University of British Columbia Vancouver, Canada

Date APRIL 28/99

Abstract

An airline caterer seeks to provide a meal quantity for each flight that closely matches final on-board passenger load. Faced with preparation lead-time, the caterer must estimate required meal quantities well in advance of departure. Passenger load may vary considerably during this lead-time, thus, adjustments are often required as more information becomes available.

In this thesis, we model the meal ordering processes at Canadian Airlines as a finite-horizon Markov decision process. The model generates policies that show the caterer how to adjust meal quantities at each decision point to minimize ordering costs. We evaluate the performance achieved with the optimal policies by applying them to a holdout dataset, and compare the results to those observed in actual practice. Next, we use the model to observe the multi-objective problem of excess meal provisioning and short meal provisioning, and estimate the cost of achieving high service levels.

This thesis finds that current practice in meal ordering at Canadian Airlines often achieves performance close to that achieved with the minimum cost policies. Application of the model to a group of 40 flights yielded estimated savings of \$4,500 per month, while reducing the number of short catered flights by 42%.

;

Abstract	ii
List of Figures	iv
List of Tables	v
Acknowledgement	vi
I. Introduction	1
A. Inflight Catering	1
B. Markov Decision Processes	1
C. Meal Ordering as a Markov Decision Process	2
D. Application to Canadian Airlines	3
II. Background	5
A. Problem Description: Why is meal ordering so difficult?	5
B. Current Meal Ordering Performance	6
C. The Meal Ordering Process	8
D. Other Technical Options	. 10
III. Model formulation	.11
A. Terminology and Assumptions	.11
B. A Markov Decision Process Model	.12
C. Modeling Passenger Demand	.14
D. Performance Measures	.18
IV. Model Development	. 19
A. Aggregate Model	.20
B. Generating Transition Probability Matrices	.20
C. Specifying Model Parameters	.23
D. Tie-Conditions and Rounding Error	.24
E. Model Results	.24
V. Results and Sensitivity Analysis	.26
A. Basic Results	.26
B. Overage versus Shortage	.31
C. Comparison of Current Performance to Model Performance	
D. Sensitivity of Performance to Probability Density Function Specifications	
VI. Conclusion	.42
Bibliography	
Appendix A. Overage versus Shortage Curves	44
Appendix B. Effect of Passenger Demand on Subsequent Flights to the Same Destination	.45
Appendix C. Modeling Differences in Passenger Load Between Pre-Departure Time Intervals with	40
Seasonality	48
Appendix D. Markov Decision Process Implementation	

List of Figures

.

Figure 1. Example sample paths of pre-departure passenger load	5
Figure 2. Distribution of provisioning error	6
Figure 3. Meal quantity versus passenger load	7
Figure 4. Distribution of meal quantity by passenger load	7
Figure 5. Distribution of passenger load	8
Figure 6. Timeline of events	9
Figure 7. Distribution of differences between successive pre-departure intervals	15
Figure 8. Scatterplot of difference versus passenger load	16
Figure 9. Boxplots of differences versus passenger load	16
Figure 10. Generating a transition probability matrix	17
Figure 11. Program structure	19
Figure 12. Empirical probability density function	21
Figure 13. Normal probability density function	22
Figure 14. Modeled probability density function	22
Figure 15. Methodology for generating transition probability matrices	23
Figure 16. Example ordering policy	25
Figure 17. Sample meal ordering policies	27
Figure 18. Distribution of provisioning error	28
Figure 19. Distribution of provisioning error: model versus actual	29
Figure 20. Distribution of provisioning error over increasing terminal cost	30
Figure 21. Example average overage versus shortage curve	31
Figure 22. Possible model results	32
Figure 23. Overage versus shortage curves	33
Figure 24. Distribution of proportion of flights short-catered	35
Figure 25. Additional average overage versus actual proportion of flights short-catered	37
Figure 26. Additional average overage versus actual average overage where shortage does not occur	37
Figure 27. Distribution of actual provisioning error	38
Figure 28. Methodology for determining best performance	40
Figure A-1. Overage versus shortage with varying levels of η	44
Figure A-2. Scenario 1: passenger load shifts to the later flight	45
Figure A-3. Scenario 2: passenger load shifts to the earlier flight	46
Figure A-4. Distribution of R^2 improvement from day of week indicators	50
Figure A-5. Distribution of R^2 improvement from seasonal indicators	52
Figure A-6. Distribution of R ² improvement from day of week and seasonal indicators	52
Figure A-7. Dimensionality of subsequent costs matrix	53
Figure A-8. Dimensionality of all possible immediate and subsequent cost combinations	54

List of Tables

.

Table 2. Distribution of provisioning error: model versus actual28Table 3. Performance indicators: model versus actual29Table 4. Bin sizes used in sample33Table 5. Model results by flight group34Table 6. Impact of the optimal policies on overage costs34Table 7. Impact of the optimal policies on number of flights short-catered35Table 8. Costs associated with achieving a zero shortage service level36Table 9. Monthly estimates of additional costs39Table 10. Effects of varying levels of alpha and bin size: Flight 1.0641Table 11. Effects of varying levels of alpha and bin size: Flight 3.0741Table A-1. Correlation of passenger load differences between subsequent flights47Table A-2. Describing seasonality through regression coefficients51	Table 1. Timing of cost components	13
Table 3. Performance indicators: model versus actual.29Table 4. Bin sizes used in sample33Table 5. Model results by flight group.34Table 6. Impact of the optimal policies on overage costs34Table 7. Impact of the optimal policies on number of flights short-catered35Table 8. Costs associated with achieving a zero shortage service level36Table 9. Monthly estimates of additional costs39Table 10. Effects of varying levels of alpha and bin size: Flight 1.0641Table 11. Effects of varying levels of alpha and bin size: Flight 3.0741Table A-1. Correlation of passenger load differences between subsequent flights47Table A-2. Describing seasonality through regression coefficients51	Table 2. Distribution of provisioning error: model versus actual	28
Table 4. Bin sizes used in sample33Table 5. Model results by flight group34Table 6. Impact of the optimal policies on overage costs34Table 7. Impact of the optimal policies on number of flights short-catered35Table 8. Costs associated with achieving a zero shortage service level36Table 9. Monthly estimates of additional costs39Table 10. Effects of varying levels of alpha and bin size: Flight 1.0641Table 11. Effects of varying levels of alpha and bin size: Flight 3.0741Table A-1. Correlation of passenger load differences between subsequent flights47Table A-2. Describing seasonality through regression coefficients49Table A-3. Describing seasonality through regression coefficients51	Table 3. Performance indicators: model versus actual	29
Table 5. Model results by flight group	Table 4. Bin sizes used in sample	
Table 6. Impact of the optimal policies on overage costs34Table 7. Impact of the optimal policies on number of flights short-catered35Table 8. Costs associated with achieving a zero shortage service level36Table 9. Monthly estimates of additional costs39Table 10. Effects of varying levels of alpha and bin size: Flight 1.0641Table 11. Effects of varying levels of alpha and bin size: Flight 3.0741Table A-1. Correlation of passenger load differences between subsequent flights47Table A-2. Describing seasonality through regression coefficients49Table A-3. Describing seasonality through regression coefficients51	Table 5. Model results by flight group	34
Table 7. Impact of the optimal policies on number of flights short-catered35Table 8. Costs associated with achieving a zero shortage service level36Table 9. Monthly estimates of additional costs39Table 10. Effects of varying levels of alpha and bin size: Flight 1.0641Table 11. Effects of varying levels of alpha and bin size: Flight 3.0741Table A-1. Correlation of passenger load differences between subsequent flights47Table A-2. Describing seasonality through regression coefficients49Table A-3. Describing seasonality through regression coefficients51	Table 6. Impact of the optimal policies on overage costs	34
Table 8. Costs associated with achieving a zero shortage service level36Table 9. Monthly estimates of additional costs39Table 10. Effects of varying levels of alpha and bin size: Flight 1.0641Table 11. Effects of varying levels of alpha and bin size: Flight 3.0741Table A-1. Correlation of passenger load differences between subsequent flights47Table A-2. Describing seasonality through regression coefficients49Table A-3. Describing seasonality through regression coefficients51	Table 7. Impact of the optimal policies on number of flights short-catered	35
Table 9. Monthly estimates of additional costs.39Table 10. Effects of varying levels of alpha and bin size: Flight 1.06.41Table 11. Effects of varying levels of alpha and bin size: Flight 3.07.41Table A-1. Correlation of passenger load differences between subsequent flights47Table A-2. Describing seasonality through regression coefficients.49Table A-3. Describing seasonality through regression coefficients.51	Table 8. Costs associated with achieving a zero shortage service level	36
Table 10. Effects of varying levels of alpha and bin size: Flight 1.06	Table 9. Monthly estimates of additional costs	39
Table 11. Effects of varying levels of alpha and bin size: Flight 3.07	Table 10. Effects of varying levels of alpha and bin size: Flight 1.06	41
Table A-1. Correlation of passenger load differences between subsequent flights47Table A-2. Describing seasonality through regression coefficients49Table A-3. Describing seasonality through regression coefficients51	Table 11. Effects of varying levels of alpha and bin size: Flight 3.07	41
Table A-2. Describing seasonality through regression coefficients	Table A-1. Correlation of passenger load differences between subsequent flights	47
Table A-3. Describing seasonality through regression coefficients	Table A-2. Describing seasonality through regression coefficients	49
	Table A-3. Describing seasonality through regression coefficients	51

Acknowledgement

The completion of this thesis was made possible by the encouragement and assistance of a number of people.

I would like to express my sincere appreciation to my thesis supervisor, Professor Martin L. Puterman, for all of his encouragement, advice, and creativity. His enthusiasm and support were always extremely helpful.

I would also like to acknowledge Acting Dean Derek Atkins, and Professor Nancy Heckman for their time and input as my thesis committee. In addition, I extend sincere thanks to Dr. Mark E. Lewis, for his technical input and help in refining the content of this thesis.

Special thanks go to the Centre for Operations Excellence (COE), for providing a forum for applied research projects and providing financial assistance. I would like to thank the Assistant Director, Stephen Jones for his keen insight early in the project, and subsequent input and feedback. Also, many COE students have offered assistance and input along the way, and I appreciate their contribution.

I offer many thanks to Canadian Airlines for sharing this challenging problem with the COE, and for providing financial assistance. Specifically, I would like to thank the Vice President of Inflight Services Marshall Wilmot, Director of Catering Services, Michael Joss, Manager Aircraft Provisioning, Manly Sitter, and Huw Harris former manager of Aircraft Provisioning. I would also like to thank Spencer Dayne and Mandy Zhou from the Inflight Services finance department, and David Lee and Don Buchanan from LSG Lufthansa Skychefs. All staff mentioned were instrumental in providing resources, describing the systems, and facilitating access to data.

I would like to express sincere gratitude to my family, for their encouragement and support despite the distance. Thanks to Lee Geyer for inspiring me to pursue a higher education and for keeping me busy along the way. Also I would like to thank my good friend Guy Rispoli, for consistently adding client value, and Larissa Petrillo, for her editorial contributions and all her support.

Financial assistance received from the Natural Science and Engineering Research Council (NSERC) through the Industrial Postgraduate Scholarship was greatly appreciated.

I. INTRODUCTION

The passenger airline industry operates on low profit margins with many competitors. Airline carriers sustain profitability through operational efficiency improvements and by maintaining or increasing market share.

Classic applications of operations research in the airline industry include yield management, route selection, and optimization of flight and maintenance schedules. Other applications include crew scheduling, analysis of customer demand, and fuel consumption control. The majority of these applications focus on operations central to the airline.

In the current competitive environment some carriers are attempting to generate savings through efficiency improvements in their periphery operations. Savings generated from increased efficiency may be directed toward improving customer service. Inflight meal provisioning is one such area worthy of pursuit as it involves high volumes and significant costs, and has direct impact on customer service.

A. Inflight Catering

Inflight meal catering refers to the provisioning of a meal service for each passenger during a flight. This service is typical of long duration flights. The complexity of the meal service varies by class of passenger service and flight destination.

Excess costs are incurred when the meal quantity on the aircraft exceeds the passenger load. Furthermore, customer service costs are incurred when the meal quantity is lower than the passenger load. The caterer seeks to provide a meal quantity that closely matches the passenger load at departure.

At several key decision points, prior to the departure of a flight, the final passenger load is estimated and the meal order is adjusted accordingly. Typically, the caterer may access information concerning tickets sold, passengers checked in, and the number of stand-by passengers. The final passenger load estimate is based on this information and the personal judgement of the catering staff. The estimation, monitoring and adjustment of the required meal quantity is referred to as *meal ordering*.

B. Markov Decision Processes

In a sequential decision problem, a system is described by the different states it may hold. The decision maker at each sequential time stage chooses an action that will cause the system to either change from one state to another, or to stay the same. The change may occur deterministically or probabilistically, and typically there are costs or rewards associated with the actions. A Markov decision process is defined as a "stochastic sequential decision problem in which the set of actions, the rewards, and the transition probabilities depend only on the current state of the system and the current action selected; the history of the problem has no effect on current decisions" (Puterman, 1992).

The time points at which decisions are made define whether the problem is discrete-time or continuoustime, and whether the time horizon is finite or infinite. The formulations we are primarily concerned with in this thesis are discrete-time finite horizon problems, where the set of time points consist of $T = \{N, N - 1, ..., 0\}$. Note that the stages are numbered in decreasing order, indicating the number of time points remaining until the terminal stage. In the finite horizon problem, terminal costs incurred in the final stage (t=0) effect the decisions made in the stages preceding it.

In the one-period problem, the system is initially in a state s_1 , and the decision maker is faced with a choice of actions. Depending on the current system state s_1 and the action a_1 chosen, the system will be in

state s_0 at the terminal stage, with a probability $Pr(s_0 | s_1, a_1)$ and immediately following the decision an immediate reward or cost r_1 is incurred. A value v_0 is associated with occupying a specific state in the terminal stage. An expectation of the value $E(v_0)$ is based on the action chosen, the transition probabilities, and the terminal rewards associated with occupying accessible states. The optimal decision rule chooses the action that maximizes the immediate reward and expected subsequent value. A value function v_1 is the sum of r_1 and $E(v_0)$, and provides the value of occupying a specific state and choosing a specific action. Here, there was only one period, thus, the decision maker is concerned with choosing the single best action.

The two-period problem is more complex. The decision maker must choose an action at time t=2 and, as a result of that action, the system will be in state s_1 at time 1, with a probability $Pr(s_1 | s_2, a_2)$ incurring an immediate reward r_2 . The decision maker must then repeat the one-period problem, with the system in state s_1 . Again, the optimal decision rule chooses the action that maximizes the immediate reward r_2 and the expected subsequent value $E(v_1)$. The value function here is denoted by v_2 . Note that the decision maker seeks the sequential set of decisions that maximizes the total value $v_2 + v_1$ over the entire decisionmaking horizon.

In the N-period problem, the decision maker chooses an action at time t=N, and as a result of that action the system will be in state s_{N-1} at time N-1, with a probability $Pr(s_{N-1} | s_N, a_N)$ and an immediate reward r_N . The probability $Pr(s_{N-1} | s_N, a_N)$ is called the transition probability function. The sequential set of decisions that maximize total value are obtained by evaluating the one-period, two-period, ..., N-1 period problems, together with the immediate reward r_N . At each stage, the maximum value is obtained by choosing the action a_N that maximizes immediate reward and expected value from the remaining actions.

Maximum Value_N = Immediate Reward_N + Expected Value from Remaining Actions_{N-1} Maximum Value_N = $r(s_N, a_N) + E(v_{N-1})$

The problem must be solved using backward induction where the results of the one-period problem are required to solve the two-period problem, and the results of the two-period problem are required to solve the three-period problem, continuing until a solution to the N-period problem is obtained.

A function that specifies an action given a state s and time period t, is called a *decision rule*. Note that the decision rule only depends on the current state of the system. A collection of sequential decision rules describing the actions to be used over the entire problem horizon is referred to as a *policy*.

C. Meal Ordering as a Markov Decision Process

The meal ordering decision making process may be modeled as a finite horizon Markov decision process. The components and assumptions of the model are consistent with the real-life application. The states of the system may be represented by the number of passengers and the number of prepared meals at a given decision epoch. The reward function is preferably based on actual corporate prices or costs of meal provisioning. Generally speaking, the costs may include basic meal costs, transportation costs, return costs, excess meal costs or short meal costs. Each of these costs may be fixed, linear, nonlinear or a combination of these. The transition probabilities represent the probability of the passenger load increasing or decreasing a certain amount between two sequential decision points. We base these probabilities on models of historical passenger demand.

Such a model will generate optimal cost meal ordering policies. These policies will specify the optimal meal order quantity adjustment given the current booked passenger load, the current meal quantity produced, and the interval of time prior to departure. Though the model will include some states that are almost never accessed, the results provide a contingency plan for those situations should they arise. For

example, a flight will never have a full set of meals and be without passengers, but the model provides an action should that event occur. In addition, the model parameters may be adjusted to evaluate possible changes to business practice.

D. Application to Canadian Airlines

In this thesis, we model the meal ordering processes at Canadian Airlines as a finite horizon Markov decision process. Canadian Airlines is a full service international airline offering passenger and cargo air transportation services to more than 240 destinations worldwide. Canadian Airlines carried approximately 11 million passengers in 1998.

The airline has three high-volume hubs of passenger activity: Vancouver, Toronto, and Calgary. Approximately 75% of the meals for North American routes are produced in kitchen facilities at these departure stations. Preliminary analysis of meal ordering performance identified that costs associated with excess meal provisioning are approximately \$1.8 million annually. The costs associated with under-provisioning flights have not been estimated, as the considerations are far more complex and depend on corporate marketing strategies.

Though the model may be applied to all flight routes, the analysis was primarily conducted with a group of flights departing from a single high activity station. We further restrict the analysis to consider only the economy class and not the business class service. The model ignores passenger meal type selection such as chicken versus beef.

While no directly related research has been conducted in the area of optimal policies for airline meal ordering some components of the problem are similar to other applications. In meal ordering we seek to provide a meal order that closely matches final passenger load. This may be thought of as an inventory control or news vendor problem, where the goods are perishable.

In the news vendor problem, the decision-maker attempts to select the order quantity that maximizes profit given a demand function, order costs, salvage value and selling price. In the meal ordering problem the demand must be satisfied only at the terminal point, and we have several opportunities to adjust the order. Though perishable inventory theory seems applicable to this problem, the majority of the literature is concerned with an infinite horizon problem, with goods perishing over several periods (Nahmias, 1982).

Another aspect of the problem involves the prediction of final passenger load. An accurate prediction of final passenger load solves any difficulties in meal ordering, however this is a non-trivial task. Prediction of passenger load is a major component of yield management, thus much research has been conducted in this area (Smith, 1992, Sun 1998).

Large clothing retailers experience a similar problem, where the manufacture of next year's fashions requires significant lead time, and more information becomes available as time approaches the actual season. Highly fashionable clothing items have a risk of losing their value shortly after the season has passed, yet the retailer seeks to meet most demand. This problem is larger in comparison to the meal ordering problem and the estimates for demand are much more subjective.

Throughout this thesis the processes, costs and current level of service are described based on interviews with staff at Canadian Airlines, and data obtained from their systems. In accordance with an agreement on confidentiality of information, some of the detailed results are not provided, or are generalized so that private information is not shared with their competitors. The interested reader is advised to request access to more detailed results directly from Canadian Airlines.

The remaining chapters of this thesis are organized as follows. In Chapter II, we describe the problem in greater detail offering insight into the factors that make it difficult. We describe the nature of the current meal ordering performance, and provide some background on the involved processes and costs.

In Chapter III we formulate the meal ordering process as a Markov decision process. We describe the states, actions, transition probabilities, decision epochs and costs as they apply to meal ordering. We describe the general methodology used to build the transition probability matrices, and introduce the aggregation of model state space.

Details on the model development are provided in Chapter IV with a discussion of the coding, the required parameters, and output generated. Other implementation issues such as tie-condition optimality and round-off error are also discussed. In addition, refinements to the transition probability matrices are described.

In Chapter V model results are provided with a sensitivity analysis on the costs and probabilities. We begin with a discussion of basic model results and extend the analysis to the overage versus shortage problem. Cost savings are provided where the model performs closest to current practices, as are the additional costs required to achieve a high level of service. We finish the chapter by performing a sensitivity analysis on some of the model parameters that affect the transition probability matrices.

4

II. BACKGROUND

A. Problem Description: Why is meal ordering so difficult?

Inflight meal catering involves the preparation, assembly, and delivery of a quantity of meals for a given flight. Due to the large volumes, the entire process may require several hours of lead time. Within this lead time the booked passenger load may fluctuate substantially due to late ticket purchases, cancellations, and passenger upgrades from economy to business class. The following plot shows how the passenger load fluctuates at various points in time prior to the departure of a flight. Each line represents a single flight.



Figure 1. Example sample paths of pre-departure passenger load

These series represent eight consecutive instances of the same flight departing on a single day of the week. We can see that the passenger load does not vary over the pre-departure intervals with an obvious pattern, except for the final transition where passenger load decreases. Note that the passenger load exceeds the capacity at some pre-departure time intervals due to overbooking. Throughout this time period, the caterer estimates the final passenger load based on the available information and adjusts the meal order at key decision points.

Canadian Airlines provides two classes of service: economy and business class. Passengers upgrading from economy to business class with more than 24 hours notice are provided a business class meal service. Last minute upgrades are advised that there may not be a meal available for them. These last minute upgrades create an excess meal in the economy section and a shortage in the business section. Serving an economy class meal in the business class cabin is not allowable, as this practice is believed to degrade the business class experience for other passengers.

The Canadian Airlines fleet includes five different aircraft models. Each aircraft model has a unique galley configuration that holds trolleys and carriers for meal provisioning. Differences in galley configuration make it difficult to move excess meals from one flight to another. In addition, Canadian

5

Airlines utilizes a full equipment policy. Under this policy, each flight carries a full set of trays, dishes and cutlery, eliminating the need to periodically adjust equipment inventory between kitchen sites. This practice, however, makes it difficult to adjust the meal order after it has been assembled.

Meal provisioning involves further constraints from health, production and corporate policies, which are described further in the section on the *Meal Ordering Process*. The considerations presented above help illustrate the difficulty in meal provisioning within a tight performance boundary.

B. Current Meal Ordering Performance

In meal ordering, a flight is described as *even-catered* when the meal quantity matches the final passenger load, *over-catered* when the meal quantity exceeds the final passenger load, or *short-catered* when the meal quantity is below the final passenger load. *Overage* and *shortage* refer to the quantity of meals over-catered and short-catered respectively. Overage and shortage are measures of *provisioning error*, which is meal quantity less final passenger load.

Meal ordering performance may be described with a variety of measures:

- proportion of flights short-catered, over-catered, or even-catered
- total overage/shortage over a period of time
- overage/shortage expressed as a percentage of final passenger load
- number of flights with overage exceeding 5% of the final passenger load
- overage associated with flights where overage exceeds 5% of the final passenger load
- distribution of provisioning error by flight

All performance measures are concerned with frequency or magnitude of errors in meal provisioning. The distribution of provisioning error on a flight-by-flight basis provides insight into the magnitude of the problem. Figure 2 shows the distribution of provisioning error for single-stage flights departing from one station with a 108 seat capacity aircraft, observed over a six month period.



Figure 2. Distribution of provisioning error

While performance varies by aircraft type, day of the week, season and flight route, these variables only explain a marginal proportion of the system-wide variance observed. The variability in meal ordering performance for this group of flights is reasonably large, with a standard deviation of 7.87 meals. This variability is also apparent in the following plots of meal quantity versus passenger load. The spike that appears where provisioning error is 0 reflects the flights that were even-catered.

Figure 3 depicts the quantity of meals versus the final passenger load. Each point represents a single instance of a flight. Observations above the 45° line represent flights that were over-catered. Observations below the 45° line represent flights that were short-catered. Flights that were even-catered will fall exactly on the 45° line. The top right shows truncation of variability in meal ordering performance due to capacity of the aircraft.



Figure 3. Meal quantity versus passenger load

The density of the observations in the previous plot is not immediately obvious. In Figure 4 the same data is presented, however the distribution of meal quantities by passenger load are illustrated using boxplots. Flights with passenger load close to the capacity of the aircraft are shown to be frequently even-catered.



Figure 4. Distribution of meal quantity by passenger load

Finally, the distribution of flights by passenger load is of interest. Figure 5 illustrates that a large proportion of flights depart at full capacity. Thus, the distribution of passenger load and the associated meal ordering performance explain the peak observed in the distribution of provisioning error.



Figure 5. Distribution of passenger load

The above figures illustrate the current variability in meal ordering performance. The distribution of provisioning error indicates that there is a tendency to over-provision these flights. Approximately 39% of the flights shown have overage exceeding 5 meals. The Canadian Airlines standard for acceptable overage is 4.2% of the passenger load.

Approximately 7.5% of the flights departed with shortage exceeding 5 meals. Canadian Airlines has not yet established an acceptable level of shortage. Some considerations may include:

- How many passengers historically refused a meal on the flight?
- What is the likelihood that failure to provide a meal will result in the loss of a customer?
- What is the expected cost of losing a customer?

These considerations indicate that short-catering costs may be nonlinear with respect to shortage. There may be a high probability that at least one passenger refuses a meal on a given flight, but a very low probability that many passengers will refuse a meal. Thus, if a flight is short one meal, this may have a low short-catering cost. Alternatively, if the flight is short by many meals, there is an increased possibility that a passenger that wanted a meal did not receive one.

C. The Meal Ordering Process

Meal ordering consists of two main stages: a production process and an adjustment process. The production process primarily takes place in the flight kitchen and involves the preparation, assembly and delivery of a quantity of meals for a given flight. The adjustment process involves alterations to a meal order after it has left the kitchen. These processes are described for a single flight for the sake of simplicity.

A meal ordering clerk is responsible for the production process. The clerk schedules resources, plans food production, and coordinates staff to build a meal order that closely matches the meal requirement for a given flight. The process spans 24 hours with responsibilities passed from shift to shift. At key points in the production process, the clerk must estimate the meal quantity requirement for the flight. The clerk has access to three sources of information. One source shows the number of tickets booked for the flight. Another source provides the clerk with a forecast of the passenger load based on the flight history. The third source of information is a forecast of the required meal quantity. The meal ordering clerk estimates the meal requirement using these sources combined with personal judgement.

The production process begins with setting a food production and staffing schedule for a flight departing on the following day. Depending on the type of meal, some foods are prepared and cooked as early as 12 hours before the flight departs. The meal requirement is estimated prior to setting the schedules and re-estimated at the point of food preparation.

A standard inflight meal consists of a tray with a casserole dish, a salad item, a dessert item, cups and cutlery. A casserole typically consists of a starch, a vegetable, and a meat/poultry/fish item. These components are cooked and chilled immediately. In the assembly stage, chilled components are assembled into casseroles that are stored in oven carriers. The tray set-up refers to the assembly of all items, except for the casserole, onto the tray. Assembled trays are stored in larger carriers or trolleys. Additional carriers are assembled for beverage service, business class presentation trays and commisary items. The meal order is assembled so that the number and size of the carriers exactly match the aircraft galley configuration. The assembly stage begins at approximately six hours prior to departure based on a revised meal requirement estimate.

Once the entire meal order is assembled into carriers, it is moved to a holding fridge. At three hours prior to departure the meal ordering clerk reviews the meal requirement estimate and adjusts the meal order as required. This is the last opportunity to adjust the meal order at the kitchen. After the order adjustment is complete, staff verify that the entire order is correct. After the meal order has been checked it may be delivered to the aircraft. Depending on the flight, the actual delivery may occur between three hours prior to departure, and a half hour prior to departure.

The adjustment process involves alterations to the meal order after it has been delivered to the aircraft. Adjustments are made through a small capacity van, which travels between the flight kitchen and the aircrafts. The van driver maintains radio contact with the kitchen and is alerted when an aircraft requires an adjustment. The driver verifies the adjustment with inflight staff and physically carries the meals onboard the aircraft. A flight may need an adjustment if the passenger load has changed substantially after the meal order was delivered.

Timeline of events

The following depicts the timeline of events that are pertinent to modeling the catering decision process.





The key decision points for the production process occur at eighteen, six, and three hours prior to departure. Subsequent decisions are made in the adjustment process in the remaining pre-departure time.

Some flights have a short turnaround time where the aircraft lands and departs in less than 3 hours. In these cases the meal order remains in the holding fridge until order delivery, thus, some late meal order adjustments may not necessarily be made with the van.

Operating Costs and Penalties

A simplified description of the meal provisioning cost structure is as follows. The cost of a meal service varies with the class of service and menu. These costs include all raw food and labor costs that are directly involved in preparing and assembling the meal order. Meal costs increase with meal quantity.

An order adjustment that occurs within the production process does not incur a penalty cost, whereas adjustments made with the van incur a 50% penalty cost plus a fixed van delivery charge. Note that a van has a capacity of 30 meals.

Costs of overage are estimated as the cost of the excess meals. Costs associated with shortage are estimated as the expected loss in customer goodwill. These costs may be based on the likelihood of losing a customer given that a meal service was not provided, where lost revenue represents the cost.

The costs associated with the main meal order delivery, and dishwashing are fixed.

D. Other Technical Options

In addition to a Markov decision process approach, some other options are available for improving meal ordering performance. Improved accuracy in estimating the meal quantity requirement may reduce the need for an adjustment process. Such forecasting may involve using passenger specific data to estimate patterns in passengers switching between flights. Improved forecasts will also improve the performance of the Markov decision process approach.

Alternatively, the systems and processes may be modified to streamline the adjustment process. The meal quantity may be made such that the order is always adjusted with an increase at the last adjustment opportunity. This would require a "meal bank" solution whereby the meals used for adjustment would be sufficiently generic so they may be used with different flights. The means for making the order adjustment would have to be simplified so that it may be accomplished quickly and reliably.

The current objective in meal provisioning is to provide a meal quantity that closely matches the final passenger load. A preferable meal quantity would be the actual number of passengers on board who wanted a meal. Analysis of actual meal consumption patterns may identify routes that should be intentionally short-catered. Obtaining this data would involve recording the number of meals consumed on a flight by flight basis, as it is not currently captured by any Canadian Airlines systems.

III. MODEL FORMULATION

A Markov decision process produces an optimal ordering policy when considering ordering costs. Terminology, assumptions, and formulation of the model are described, followed by a discussion of how the passenger demand is modeled.

A. Terminology and Assumptions

Some terms were introduced in the previous section on the meal ordering process. Some of these terms are known within the airline industry but are not standard within academic literature. To avoid confusion, and to be consistent with existing literature on the subject, the terms are described below.

Pre-departure time	The length of time remaining before the aircraft departs.
Decision Epoch	The time at which an adjustment to the meal order can be made.
Passenger load	The number of booked and stand-by passengers. At the point of departure this is the number of on-board passengers.
Meal quantity	The number of meals to be prepared and assembled for a given flight. Depending on the pre-departure time, this may be in production, in the holding fridge, or onboard the aircraft.
Capacity	The passenger capacity of the aircraft.

In order to model the meal ordering process some assumptions regarding the process, the passenger demand, and the costs are required. These assumptions and their associated implications are described.

- 1) The passenger demand on one flight does not affect subsequent flights. Transitions in passenger demand between the pre-departure decision epochs are modeled independently on a flight-by-flight basis, thus any information regarding other flights is not considered. This assumption would be valid if the following holds. Two flights, A and B, with the same origin-destination pair depart on the same day at different times. If flight A departed first and was excessively overbooked, the excess passengers do not affect the passenger load on flight B.
- 2) While absolute passenger load varies by season and day of the week, the passenger demand for these flights may be modeled together. We do not seek to predict absolute passenger load, but are attempting to describe how the passenger load will change between the pre-departure intervals.
- 3) Exceptional changes in passenger load due to special events, holidays or aircraft maintenance problems will be removed from the modeling of the transition probabilities. Due to the large amount of variability these effects are not always obvious.
- 4) During the adjustment process, when the meal order is already at the aircraft, order adjustments may be made at two hours and one hour prior to departure.
- 5) Aggregate models may be used for modeling transitions in passenger load between pre-departure intervals. In an aggregate model, a single seat may represent a batch of seats (2, 3, 6, etc.). This practice becomes necessary when there are a small number of observations available for modeling the passenger demand.

- 6) Meal ordering involves the following variable costs: raw meal costs, a return penalty, van delivery charge and customer goodwill cost when a customer does not receive a meal due to short-catering. Fixed costs such as one-time transportation and dishwashing are not considered in the model.
- 7) There is a fictitious increasing cost of producing and assembling the meal order as the departure time draws near. This cost may be interpreted as the expected cost of delaying the flight departure time due to late meal provisioning.
- 8) All meals ordered in a pre-departure time interval are ready and delivered in the following time interval.
- 9) The meal order quantity can never exceed the capacity of the aircraft.

For the group of flights analyzed, assumptions 1 and 2 appear to be true. The analysis surrounding these conclusions is provided in the appendices. Assumptions 4, 6, 7, 8, and 9 are based on the current processes and practices in place. The remaining assumptions 3 and 5 are required for the basic operation of the model. The validity of these assumptions should be checked prior to applying the model to other routes.

B. A Markov Decision Process Model

The following notation is used in describing the model:

 $mq_t = meal quantity at time t$ $pl_t = passenger load at time t$ a = meal quantity after adjustment

The state space consists of all possible values for the passenger load and meal quantity pair (pl, mq). Values for each span from zero to the capacity of the aircraft.

$$S = \{0, \dots, Capacity\} \times \{0, \dots, Capacity\}$$

The decision epochs are $T = \{0,...,5\}$ where

t	Time before departure
0	Departure time
1	1 hour
2	2 hours
3	3 hours
4	6 hours
5	36 hours

Note that the range of possible adjustments to the meal order are limited by the current state. For example, if the current meal quantity is 12 meals, then up to 12 meals may be removed from the aircraft and up to capacity minus 12 meals may be added. The set of state dependent actions are the meal order quantity after adjustment.

$$a_{meal + order \atop order_t adjustment_t} = mq_{t-1} \in A = \{0, ..., Capacity\}$$

The transition probabilities represent the probability that the passenger load will be pl_{t-1} at decision epoch t-1, given that the passenger load was pl_t at decision epoch t. Note that the probabilities are independent of the action, and the current meal quantity. We denote these by $Pr(pl_{t-1} | pl_t)$. These are assembled into a transition probability matrix. This square matrix has dimensions $\{Capacity + 1\} \times \{Capacity + 1\}$, with the $(pl_t, pl_{t-1})^{\text{th}}$ component containing $Pr(pl_{t-1} | pl_t)$.

The cost/reward function is as follows:

$$r_t((mq_t, pl_t), a) = raw action costs + return penalty_t + van delivery charge_t + terminal costs + late penalty_t$$

Notice that some of the components of the cost/reward function are time dependent. As shown in Table 1, some of these costs only apply to specific time intervals. These are indicated with an "X" at the appropriate decision epoch.

`t	Raw action costs	Return penalty	Van delivery charge	Terminal costs	Late penalty
0			· · ·	X	
1	Х	Х	Х		Х
2	Х	X	Х		Х
3	Х				Х
4	Х				Х
5	Х				

Table 1. Timing of cost components

Raw action costs refer to the basic cost of a meal. As described previously, order decreases (meals that are returned) that occur after the order has been delivered to the aircraft incur a penalty cost of 50% of the meal value. For example, a meal that costs \$10 may be returned at no penalty if the order has not been delivered to the aircraft. However, once the meal order is delivered to the aircraft a \$5 penalty fee is incurred on any order decrease. Any positive increase in the meal order quantity during the adjustment stage will result in a fixed van delivery charge.

A terminal cost is assigned to overage or shortage in the final meal order quantity. Overage costs are simply the raw costs of the excess meals. Shortage costs represent the potential loss in customer goodwill. Estimates for actual shortage costs are obtained by evaluating the model at various shortage costs, and observing the value which results in the closest match to actual performance. This will be described in more detail later in the sensitivity analysis.

The late penalty is a fictitious cost that represents the caterers' preference for building the meal order earlier than later. The flight kitchens at Vancouver, Toronto, and Calgary provision many flights every day and cannot reliably build all meal orders at the absolute last opportunity. The late penalty is small and increases as time approaches the departure time.

A Markov decision process model generates optimal policies. A policy is a collection of decision rules, one for each decision epoch, which describe how to adjust the meal order given a current meal quantity and passenger load.

We assume at decision epoch 5 the meal quantity mq_5 is zero. For a given meal ordering policy $\pi = (d_5, d_4, d_3, d_2, d_1)$, the expected total provisioning and penalty cost is:

$$v^{\pi}(0, pl_5) = E\left\{\sum_{t=1}^{5} r_t((mq_t, pl_t), d_t(mq_t, pl_t)) + r_0(mq_0, pl_0)\right\}$$

The optimal cost policy and associated costs are obtained through the following recursive value function:

$$t = 5, ..., 1$$

$$v_{t}^{*}((mq_{t}, pl_{t}), mq_{t-1}) = \min_{mq_{t-1} \in A} \left\{ r_{t}((mq_{t}, pl_{t}), mq_{t-1}) + \sum_{pl_{t-1} \in [0, Capacity]} \Pr(pl_{t-1} | pl_{t}) v_{t-1}^{*}(mq_{t-1}, pl_{t-1}) \right\}$$

$$t = 0$$

$$v_{0}^{*}(mq_{t}, pl_{t}) = r_{0}(mq_{t}, pl_{t})$$

$$mq_{t-1}^{*} \in \arg \min_{mq_{t-1} \in A} \left\{ r_{t}((mq_{t}, pl_{t}), mq_{t-1}) + \sum_{pl_{t-1} \in [0, Capacity]} \Pr(pl_{t-1} | pl_{t}) v_{t-1}^{*}(mq_{t-1}, pl_{t-1}) \right\}$$

Note that due to the possibility of multiple minima, the minimum cost action may not be unique. This will be explained in more detail in the *Model Development* section.

C. Modeling Passenger Demand

Estimating a transition probability matrix requires modeling of passenger demand. The available data consists of booked passenger load and stand-by passenger load. Passenger load is the sum of the two loads. The loads are available for the following pre-departure intervals: Post departure, $\frac{1}{2}$ hour, 1 hour, 2 hours, 3 hours, 6 hours, 9 hours, 12 hours, 18 hours, 36 hours, 7 days, and 14 days. In the meal ordering process we are concerned with transitions in passenger load between pre-departure intervals 36 hours, 6 hours, 6 hours, 1 hour, and post departure.

Several options were pursued for modeling the transitions in passenger load between the intervals. Initial efforts focused on modeling differences in passenger load. Subsequent efforts focused on using absolute passenger load. Both are described below with examples.

Differences in Passenger Load

A *difference* is the change in passenger load, for a given flight, between two successive pre-departure time intervals. The distribution of differences is of interest in modeling pre-departure changes in the passenger load. Five differences are obtained for each flight instance: 36 hours to 6 hours, 6 hours to 3 hours, 3 hours to 2 hours, 2 hours to 1 hour, 1 hour to post departure.

Modeling the differences in passenger load requires a reasonably small number of observations. However, this methodology does not capture how the transition in passenger load between two intervals may relate to absolute passenger load. For example, the distribution of differences may become more variable as passenger load increases, or the mean of the differences may shift. The following distributions are shown for the five transitions. The distributions are based on the same group of flights as presented previously in previous section on *Meal Ordering Performance*.



Figure 7. Distribution of differences between successive pre-departure intervals

We see that the distributions of differences are centered close to zero, except for the last transition. The distribution of the "one hour to post departure" differences exhibits a negative mean and comparatively large variability. This negative mean may be due to last-minute ticket cancellations or passengers missing their flight. It may also be a partial result of overbooking with some booked passengers being shifted to other flights. Alternatively, passenger load may decrease as stand-by passengers attempting to leave on an earlier flight are asked to wait for their original flight.

The distributions of the differences were modeled with Poisson and normal distributions, independent of passenger load. Specifically, we are modeling the probability of a given change in passenger load denoted as $Pr(pl_{t-1} - pl_t)$. Note that the Poisson distribution consists of only positive discrete values, thus it was required that the differences be temporarily shifted to positive values for the data modeling, based on a trimmed minimum value. Both distributions provided a comparable fit. As shown in Figure 7 some of the distributions have long tails. These tails may be modeled using a combined model of a normal distribution with a uniform distribution for the tail. This modeling has not yet been attempted.

A transition probability matrix is constructed using these distributions. A unique probability matrix is required for each of the differences. In this matrix all rows have the same probability distribution, and the contents shift to the right by one cell as the row increments downwards. Where probabilities exceed the left or right edges of the matrix, the cumulative distribution is used in the boundary cell.

As the following plot shows, the distribution of differences may vary by passenger load. This phenomenom was only observed at the final passenger load transition. This provides motivation to

consider other factors in modeling differences in passenger load between successive pre-departure intervals.



Figure 8. Scatterplot of difference versus passenger load

The relative distribution at each passenger load is of interest. In Figure 9, boxplots show the distribution of differences at increasing levels of passenger load. We can see the negative relationship between passenger load at 1 hour pre-departure and the differences. The thick line crossing through the midpoints of the boxplots represents a simple regression with passenger load.



Figure 9. Boxplots of differences versus passenger load

The plot of differences versus passenger load also identifies the censoring that occurs in the final transition. Overbooking tickets for a flight is standard practice in the airline industry. The strategy maximizes the yield of the flight given that some customers will cancel at the last minute or not show up at all. As shown in Figure 9, the effect of overbooking appears as censored demand. The line appearing at capacity represents the differences that must occur in the last transition, so that the on-board passenger load does not exceed capacity.

Absolute Passenger Load

The transition probability matrix may alternatively be developed by counting transitions from one passenger load to another between two successive pre-departure intervals for a sample of flight observations. The counts are divided by the row count total to obtain a probability matrix. This methodology captures passenger load dependent transitions, however, a large number of observations are required. The matrix will likely have sparse entries, even with 100 to 200 data points, as some states are never reached (see Figure 10 b). In the event that no observations exist for a given passenger load at any decision epoch, the row corresponding to that passenger load will be missing from the transition probability matrix. Any missing row is filled with a distribution based on the differences (see Figure 10 c).



Figure 10. Generating a transition probability matrix

The data and matrix may be reduced such that a single seat represents a batch of two or three seats. The matrix may then be expanded to the original size by dividing all row elements by the batch size.

Improvements in Demand Modeling

Analysis was conducted to determine if any covariates could help model the change in passenger load between the final two intervals. This transition is the most critical in the performance of the model. The following information was analyzed with respect to final passenger load:

- Passenger load at previous intervals
- Day of the week
- Cascade effect of passengers from previous flights to the same destination
- Season
- First/last flight of the day
- Business class passenger load at previous intervals (bumping of passengers)
- Forecasted passenger load

All variables were included in a multiple regression with a stepwise selection process. It was found that most of these factors are highly correlated with the passenger load at one hour prior to departure, resulting in this one variable as explaining the majority of the variability. The variability of the errors was only slightly smaller than the variability of the raw differences. Thus, the gain from this approach is marginal.

Analyses concerning the cascade effect of passenger from previous flights to the same destination and the seasonality effect are presented in the appendices.

Other methods available for predicting final passenger load include exponential smoothing, ARIMA modeling, and demand and booking averaging. An exponentially smoothed average is based on the passenger load at the previous pre-departure intervals. Preliminary time series analysis revealed little gain in predicting final passenger load. The demand and booked average is a variable weighted average of current bookings and smoothed final bookings.

D. Performance Measures

The product of the Markov decision process model is a minimum expected cost meal order adjustment policy. This policy is generated based on a subset of available data, and the ordering policies created are applied to holdout dataset. The performance of the model is primarily measured by its ability to closely match the final passenger load. Provisioning error is obtained for the historical performance, and for the model. Recall that provisioning error is determined on a flight by flight basis and is the meal quantity minus the final passenger load. The following performance measures are obtained for both the model and actual data:

- Distribution of the provisioning error (percentiles, trimmed mean and standard deviation)
- Percentage of flights that experienced a shortage
- Percentage of flights that experienced a shortage or overage exceeding 5 meals
- Total overage, total shortage
- Average overage, average shortage

These represent a refinement on the measures presented in the previous section on *Background*. Further measures consider the number of adjustments to the meal order, number of van deliveries or a cost comparison of the meal ordering process based on the actual versus modeled performance.

IV. MODEL DEVELOPMENT

The model described in the previous chapter can be used to obtain optimal minimum cost policies for meal ordering. In this chapter we describe some details of the actual model development, and the results that it produces.

The model is developed in SAS (Statistical Analysis Software), which is a robust data management and statistical development language. SAS also includes a matrix language, flexible graphing facilities, and code generation macros. All of these functions were employed in the model development. The overall structure of the code is as shown in Figure 11.



Figure 11. Program structure

The code creates transition probability matrices and applies the Markov decision process model to determine the optimal cost ordering policies. The code then compares the performance of the optimal policy to the actual historical performance.

Two sources of data are used. The first data source consists of the passenger load at key pre-departure time intervals. The pre-departure data is split into two subsets based on a user-specified date. The earlier subset (model dataset) is used in developing transition probability matrices. An optimal policy is generated based on the cost function and the transition probability matrices. The optimal policy is applied to the second subset (test dataset) generating a "model" meal quantity.

The second data source consists of historical records of actual final meal order quantity. This data is merged to the model meal quantities and the performance of each are evaluated based on provisioning error.

A. Aggregate Model

The number of operations performed in the model increases polynomially with aircraft capacity. Our state space consists of all possible meal quantity and passenger load pairs. Thus, as capacity increases, the state space increases quadratically and the number of actions to evaluate increases linearly. For example, a 9 seat model (10 states including the zero state) requires 1000 evaluations at each decision epoch, whereas a 19 seat model requires 8000 evaluations. Running the model with a 108 seat capacity aircraft (Airbus A320) requires approximately 10 minutes of processing time on a Pentium 200 MHz personal computer. The largest aircraft in the Canadian Airlines fleet (Boeing 747) has 380 seats in the economy class. This provides motivation to work with a simplified model while interactively developing and analyzing the model.

An aggregate model was developed in which a single seat represents a batch of seats (i.e. 2, 3, 6, etc.). The costs and the transition probabilities are scaled according to the batch size, so that the model results are realistic. To run the model at full scale, the batch size is set to a value of one.

B. Generating Transition Probability Matrices

,

As described earlier, five differences are obtained from the historical pre-departure data relating to the change in passenger load between the intervals 36 hours to 6 hours, 6 hours to 3 hours, 3 hours to 2 hours, 2 hours to 1 hour, and 1 hour to post departure. For each group, the 1st and 90th percentiles are obtained, defining the limits of a trimmed region. A trimmed mean (denoted as *Mean*_t) and standard deviation (denoted as *Standard Deviation*_t) are calculated, from which a transition probability matrix is generated. As before, pl_t refers to the passenger load at decision epoch t. Where an aggregate model is applied, the passenger load pl_t and capacity are specified in the aggregate state space.

$$\Pr(pl_{t-1} \mid pl_{t}) = \begin{cases} \Pr\{Y_{t} \ge pl_{t-1} - pl_{t}\} & pl_{t-1} = 0\\ \Pr\{Y_{t} \ge pl_{t-1} - pl_{t}\} - \Pr\{Y_{t} \ge pl_{t-1} - pl_{t} - 1\} & 0 < pl_{t-1} < Capacity\\ 1 - \Pr\{Y_{t} \ge pl_{t-1} - pl_{t} - 1\} & pl_{t-1} = Capacity \end{cases}$$

where $Y_t \sim N(Mean_t, Standard Deviation_t)$

Probability matrices are constructed where the $(pl_t, pl_{t-1})^{\text{th}}$ entry equals $\Pr(pl_{t-1} | pl_t)$. Recall that the decision epochs decrease as time progresses, and that epoch t-1 occurs after epoch t.

The transition probability matrix for the final group (differences between 1 hour and post departure), is modeled differently to capture the relationship between differences and passenger load. The differences decrease as passenger load at 1 hour increases, with the correlation ranging from -0.21 to -0.59. The relationship is modeled as a simple linear regression of the form:

 $\begin{array}{l} \textit{Difference}_{1 \ \textit{hour pre-departure}} = \beta_0 + \beta_1 p l_1 + \varepsilon \\ & \text{to post departure} \end{array}$

Where β_0 and β_1 are regression parameters, pl_1 is the passenger load at 1 hour prior to departure, and ε is the error term. The transition probability matrix is constructed as described above, using the predicted difference as the mean, and the root mean square error s_e as an approximation for the standard deviation. The root mean square error is a reasonable estimate for standard deviation where the number of observations is large and the passenger load is within the range of the data used in the regression. Preferably, the standard deviation would be based on the error term used in a prediction interval. Also, the variability of the differences are censored at high values of pl_1 , as shown previously in Figure 9. Assumptions of the linear regression model include constant variance and normality of the dependent variable. Clearly these assumptions are violated where the data is censored, however the model provides reasonable fit.

Next, a frequency table is created based on the same historical data, where the $(j,k)^{\text{th}}$ entry contains the number of flights where passenger load at interval t was j passengers, and passenger load at interval t-1 was k passengers. Each element in the table is divided by its row total to convert it into an empirical probability distribution. The distribution in some rows may be non-smooth and discontinuous due to a low number of observations. The distribution may be smoothed over neighbouring cells using a row specific distribution of differences. The modeled distribution is obtained using a weighted combination of the two probability functions.

In the following example, a given row of the frequency table consists of 7 observations: passenger load shifting from 39 passengers at interval t to 33, 35, 35, 38, 39, 39, 39 passengers at interval t-1. The empirical distribution is discontinuous and probabilities are non-zero only where observations occurred in the sample dataset.



Figure 12. Empirical probability density function

We wish to smooth the probabilities so that the distribution is continuous and extends beyond the range of the seven observations, but holds the properties of the empirical distribution. A normal distribution based on the average and standard deviation of the seven observations represents a smoothed distribution, that has lost the shape of the empirical distribution.



Figure 13. Normal probability density function

A parameter alpha is used to combine the two distributions to produce a modeled distribution. The probability for any transition is obtained by:





Figure 14. Modeled probability density function

In the example shown, the value for alpha is 0.5 giving equal weight to each distribution. This modeling is repeated for all rows in the frequency table where there are a sufficient number of observations. Note that where the number of observations in the row is equal to zero, the row will not exist in the transition probability matrix. Rows with 5 or less observations are replaced with the corresponding row in the transition probability matrix obtained using the differences.

Much effort is spent in ensuring that the transition probability matrices are as accurate as possible. We wish to model the probabilities so that they will accurately represent the transitions outside of the model dataset. Discontinuous or sparse rows in the probability matrices may result in non-optimal meal ordering policies as the model interprets certain states as inaccessible.

Figure 15 summarizes the methodology of generating the transition probability matrices.





C. Specifying Model Parameters

When running the model it is necessary to specify the following:

Flight group	Any logical statement specifying the departure station, flight number, the flight departure date range, flight departure day of the week, or time of departure.
Bin size	The aggregate level in which the model will be run. Bin size refers to the number of seats that each aggregate seat represents.
Alpha	The weight to be applied to the empirical and normal distributions to obtain the modeled distribution used in the transition probability matrices.
Hold back date	The date that splits the pre-departure data into two sets: a model dataset to generate the policies, and a test dataset to evaluate the performance of the policies.
Terminal cost	The estimated goodwill cost associated with not providing a meal service to a

passenger due to shortage, or the meal costs associated with overage.

Late cost A pre-departure time dependent cost representing a preference to not create the meal order at the absolute last opportunity.

D. Tie-Conditions and Rounding Error

In some cases, given a common state, the minimum cost is achieved by two or more different actions. Such an event is called a "tie". As described earlier, the model finds the minimum cost, given a current state and time, by evaluating the cost of increasing meal order quantities, for the following interval. The first lowest cost action observed (i.e. the action with the lowest meal order quantity) is saved as the optimal policy.

In some cases, a difference in cost between two actions is created (falsely) due to rounding error. This situation will typically arise where there is a "tie" between two or more actions. The model seeks minimum cost actions and will choose the action with the smallest cost. There are two methods for dealing with this:

- Choose a number of significant digits and scale all costs and probabilities such that the significant digits are integer values. Conduct all modeling with the scaled up figures and return the results at the original level. This method may cause an overflow error in the variable when dealing with large penalty costs.
- 2) Specify a threshold of minimum improvement. This will prevent the model from choosing actions that are \$0.00001 cheaper, as the marginal savings are negligible.

The second method was applied in the final model, as it allowed for more flexible code and has intuitive meaning.

E. Model Results

The model generates optimal ordering policies and a summary report of the model performance. The model performance is obtained by applying the optimal policies to historical data, and comparing the results to actual performance.

Optimal Policy Tables

Optimal policies are presented as five tables, one decision rule for each decision point. The tables are $\{Capacity + 1\} \times \{Capacity + 1\}$ in dimension where the (i, j)th cell contains the order adjustment given that at the decision point, the passenger load is *i* passengers and the meal order quantity is *j* meals.

A modified table is constructed where the $(i, j)^{\text{th}}$ cell contains the order adjustment plus the original meal quantity, given that at the decision point, the passenger load is *i* passengers and the meal order quantity is *j* meals. Depending on the policy, the modified table may illustrate patterns in the optimal policy better than the original table.

The cells in the table are shaded according to value. The shading allows the reader to easily see any patterns in the policy. The modified table presented in Figure 16 is an example decision rule for an 18 seat aircraft. For example, given a passenger load of three passengers, and a current meal quantity of seven meals, the decision rule presented in Figure 16 indicates that the meal quantity in the next interval should be two meals. Thus, the order adjustment at the current stage is to remove five meals.

The decision rule shown has a few bands of consistent pattern. The upper right, and lower left triangular pattern indicates a capacity effect specific to the last minute order adjustment delivery van. We can see that meal order can be adjusted up to a limit of six meals at this stage. In the main diagonal band, the ordering policy decreases the meal order by one meal. We will describe the patterns in the decision rules in more detail in the *Results and Sensitivity Analysis* section.

Meal Quantity

		0	1	2	з	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
	0	0	0	0	0	0	0	0	1	2	3	4	5	6	7	8	9	10	11	12
	1	0	0	0	0	1	0	0	1	2	3	4	5	6	7	8	9	10	11	12
	2	1	1	1	1	1	2	1	1	2	3	4	5	6	7	8	9	10	11	12
	3	2	2	2	2	2	2	3	(2	2	3	4	5	6	7	8	9	10	11	12
	4	3	3	3	3	3	3	3	¥	3	3	4	5	6	7	8	9	10	11	12
	5	4	4	4	4	4	4	4	4	5	4	4	5	6	7	8	9	10	11	12
p	6	5	5	5	5	5	5	5	5	5	6	5	5	6	7	8	9	10	11	12
03	7	6	6	6	6	6	6	6	6	6	6	7	6	6	7	8	9	10	11	12
er I	8	6	7	7	7	7	7	7	7	7	7	7	8	7	7	8	9	10	11	12
gua	9	6	7	8	8	8	8	8	8	8	8	8	8	9	8	8	9	10	11	12
SSe	10	6	7	8	9	9	9	9	9	9	9	9	9	9	10	9	9	10	11	12
Ра	11	6	7	8	9	10	10	10	10	10	10	10	10	10	10	11	10	10	11	12
	12	6	7	8	9	10	11	11	11	11	11	11	11	11	11	11	12	11	11	12
	13	6	7	8	9	10	11	12	12	12	12	12	12	12	12	12	12	13	12	12
	14	6	7	8	9	10	11	12	13	13	13	13	13	13	13	13	13	13	14	13
	15	6	7	8	9	10	11	12	13	14	14	14	14	14	14	14	14	14	14	15
	16	6	7	8	9	10	11	12	13	14	15	15	15	15	15	15	15	15	15	15
	17	6	7	8	9	10	11	12	13	14	15	16	16	16	16	16	16	16	16	16
	18	6	7	8	9	10	11	12	13	14	15	16	17	17	17	17	17	17	17	17

Figure 16. Example ordering policy

Summary Report of Model Performance

We apply the optimal policies to the holdout dataset and observe the distribution of the "model" provisioning error. We compare the "model" provisioning error to the actual provisioning error observed historically. In addition to the performance measures described in the previous *Model Formulation* section, the following output is produced after the model has run:

- A histogram of the "model" provisioning error, superimposed on a histogram of actual provisioning error historically observed.
- Pre-defined percentiles of the distribution of the provisioning error.
- The number of flights in the holdout dataset.
- The mean and standard deviation of provisioning error.
- The trimmed mean and standard deviation of the provisioning error based on observations within 1st and 99th percentiles of the original data.

The entire program is designed to evaluate various scenarios and costs. The output report allows these scenarios to be compared.

V. RESULTS AND SENSITIVITY ANALYSIS

In this section we present results of applying the model to a selected group of flights. We include a discussion of the relationship of overage and shortage, and show how the model will perform differently depending on the terminal cost value. We then present results obtained from running the model over 40 flights separately, and show for each flight the cost of achieving a pre-determined service level. Each flight represents a specific route, departing at a specific time. In the final section we briefly discuss the sensitivity of the model to changes in the probability density function for passenger demand.

A. Basic Results

The model was run using historical data for a small group of flights with the same origin-destination pair, and aircraft capacity. No grouping was made for day of the week or season. The data spans a full year between February 1998 and January 1999 inclusive. A holdback date of December 1, 1998 separates the dataset into a model dataset and a test dataset. An alpha value of 0.5 was applied, combining the empirical and normal distributions with equal weighting. The terminal cost was chosen as \$120 per short meal. The late penalty costs are \$0, \$0, \$2.50, \$2.50, and \$7.50 for the decision epochs 5 through 1 respectively. The aircraft capacity is 108 seats, and a bin size of 9 was applied, thus the aggregate model has 12 seats. In practice, a much smaller bin size would be applied. We choose a coarse bin size for the sake of simplicity.

An Optimal Meal Ordering Policy

The optimal policy presented is for an aggregate model. Thus, the 12th seat refers to the group of seats 100 to 108; the 11th seat refers to 91 to 99 and so on. The decision rules are presented in 13 by 13 tables, with the row representing the current passenger load, the column representing the current meal order quantity and the cell containing the meal order after the adjustment. All references to the passenger load and meal order quantity pertain to the aggregate model unless otherwise stated. The five decision rules that comprise the policy are presented in Figure 17.

We now describe some features of an optimal policy. The policy specifies no ordering at 36 hours predeparture (decision epoch 5), thus we see the decision rule contains all zeros. At 6 hours pre-departure (decision epoch 4), the policy specifies a meal order quantity that varies only by passenger load. The decision at this point cannot be affected by the meal order quantity because no meals could have been ordered, according to our previous decision rule. The meal order quantity is consistent over bands of passenger load. This reflects that the distribution of differences between decision epochs 5 and 4 has some pattern, or alternatively, there is some cost function that is driving this decision.

At 3 hours pre-departure (decision epoch 3), the policy makes adjustments so that the meal order quantity i) matches the passenger load when the passenger load is 5 or greater, and ii) exceeds the passenger load by one meal otherwise. This interval is the last opportunity to adjust the meal order without incurring a van delivery charge, or a return penalty cost. This helps explain the pattern of the decision rule. Note that two diagonal ridges appear, indicating that the model is reacting to a cost function or the distribution of differences.

At 2 hours and 1 hour pre-departure the meal order may be adjusted via the delivery van. A per-visit delivery charge is incurred for every order increase in these two decision epochs. The van has a fixed capacity of 4 meals, the effect of which is apparent in the final decision rule. The pattern is not as obvious at decision epoch 2, as a late penalty is involved. As the model chooses the set of decisions that minimize the expected reward over all decision epochs, it may select a meal order adjustment at decision epoch 2 that reduces the likelihood of having to incur a van delivery charge at the final decision point. Note that in the main diagonal of the final decision rule, the policy orders one more meal than required.

This may either be a result of a high terminal cost, or alternatively the distribution of differences between decision epochs 1 and 0 may have a positive mean.

Model policies for other groups of flights generally follow this form, with variations due to differences in the distribution of passenger demand. Note that these meal ordering decisions are similar to those made in practice.

(pl, mq) epoch 5

	0	1	2	3	4	5	6	7	8	9	10	11	12
0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	0	0	0	0	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0	0	0	0	0	0
6	0	0	0	0	0	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0	0	0	0	0	0
9	0	0	0	0	0	0	0	0	0	0	0	0	0
10	0	0	0	0	0	0	0	0	0	0	0	0	0
11	0	0	0	0	0	0	0	0	0	0	0	0	0
12	0	0	0	0	0	0	0	0	0	0	0	0	0

(pl, mq) epoch 4

	0	1	2	З	4	5	6	7	8	9	10	11	12
0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	1	1	1	1	1	1	1	1	1	1	1	1	1
2	2	2	2	2	2	2	2	2	2	2	2	2	2
3	3	3	3	3	3	3	3	3	3	3	3	з	3
4	5	5	5	5	5	5	5	5	5	5	5	5	5
5	7	7	7	7	7	7	7	7	7	7	7	7	7
6	7	7	7	7	7	7	7	7	7	7	7	7	7
7	7	7	7	7	7	7	7	7	7	7	7	7	7
8	7	7	7	7	7	7	7	7	7	7	7	7	7
9	10	10	10	10	10	10	10	10	10	10	10	10	10
10	10	10	10	10	10	10	10	10	10	10	10	10	10
11	10	10	10	10	10	10	10	10	10	10	10	10	10
12	10	10	10	10	10	10	10	10	10	10	10	10	10

(pl, mq) epoch 3

	0	1	2	з	4	5	6	7	8	9	10	11	12
0	1	1	1	1	2	1	1	1	1	1	1	1	1
1	2	2	2	2	2	3	2	2	2	2	2	2	2
2	з	З	3	3	3	3	4	3	3	3	3	З	3
з	4	4	4	4	4	4	4	5	4	4	4	4	4
4	5	5	5	5	5	5	5	5	6	5	5	5	5
5	5	5	5	5	5	5	5	5	6	5	5	5	5
6	6	6	6	5	6	6	6	6	6	7	6	6	6
7	7	7	7	7	6	7	7	7	7	7	8	7	7
8	8	8	8	8	8	7	8	8	8	8	8	9	8
9	9	9	9	9	9	9	8	9	9	9	9	9	10
10	10	10	10	10	10	10	10	9	10	10	10	10	10
11	11	11	11	11	11	11	11	11	10	11	11	11	11
12	12	12	12	12	12	12	12	12	12	11	12	12	12

(pl, mq) epoch 2

_	0	1	2	З	4	5	6	7	8	9	10	11	12
0	0	1	2	2	2	З	2	З	4	5	6	7	8
1	0	1	2	3	3	З	4	3	4	5	6	7	8
2	З	1	2	З	4	4	4	5	4	5	6	7	8
3	4	4	2	3	4	5	5	5	6	5	6	7	8
4	4	5	5	3	4	5	5	5	6	7	6	7	8
5	4	5	5	5	4	5	6	7	7	7	7	7	8
6	4	5	6	5	4	5	0	7	7	7	8	7	8
7	4	5	6	7	6	7	6	7	8	8	8	9	8
8	4	5	6	7	8	7	8	7	8	9	9	9	10
9	4	5	6	7	8	9	8	9	8	9	10	10	10
10	4	5	6	7	8	9	10	9	10	9	10	11	11
11	4	5	0	7	8	9	10	11	10	11	10	11	12
12	4	5	6	7	8	9	10	11	12	12	10	11	12

(pl, mq) epoch 1

	0	1	2	з	4	5	6	7	8	9	10	11	12
0	1	1	1	1	2	1	2	3	4	- 5	6	7	8
1	2	2	2	2	2	з	2	3	4	5	6	7	8
2	З	З	3	3	з	З	4	3	4	5	6	7	8
3	4	4	4	4	4	4	4	5	4	5	6	7	8
4	4	5	5	5	5	5	5	5	6	5	6	7	8
5	4	5	5	5	5	5	5	5	6	5	6	7	8
6	4	5	6	6	6	6	6	6	6	7	6	7	8
7	4	5	6	7	7	7	7	7	7	7	8	7	8
8	4	5	6	7	8	8	8	8	8	8	8	9	8
9	4	5	6	7	8	9	9	9	9	9	9	9	10
10	4	5	6	7	8	9	10	10	10	10	10	10	10
11	4	5	6	7	8	9	10	11	11	11	11	11	11
12	4	5	6	7	8	9	10	11	12	12	12	12	12

Figure 17. Sample meal ordering policies

Model Performance

The histograms shown in Figure 18 depict the distribution of provisioning error obtained by applying the optimal policies to the test dataset. Recall that provisioning error is the final meal order quantity minus the passenger load at departure. *Model* provisioning error refers to the error obtained by applying the optimal cost policies. *Actual* provision error refers to the error observed in the historical data. We see that the performance of the model is comparable to the historical performance.





The distribution of the provisioning error obtained with the optimal policies has a lower mean and tighter distribution in comparison to the actual provisioning error. The mean and standard deviation of provisioning error are provided in Table 2, and boxplots of the distributions are provided in Figure 19.

	Provisioning Error			
Description	ACTUAL	MODEL		
Number of observations	120	120		
Mean	9.81	7.99		
Standard deviation	8.46	6.96		





Figure 19. Distribution of provisioning error: model versus actual

The average overage of the flights is obtained by dividing the total overage observed in the test dataset by the total number of flights. As shown in Table 3, the average overage obtained with the optimal policies is lower than the actual average overage observed in practice (8.33 versus 10.19 meals). The optimal policies outperform the actual practice in terms of overage and shortage exceeding 5 meals.

Description	ACTUAL	MODEL
Average overage	10.19	8.33
Proportion of flights with overage exceeding 5 meals	62.5%	55.8%
Proportion of flights with shortage exceeding 5 meals	1.7%	0.8%
Proportion of flights with any shortage	6.7%	6.7%

Table 3. Performance indicators: model versus actual

In this example, the optimal policies result in meal ordering that outperforms actual practice. Given another selection of bin size, alpha, late penalty costs, and terminal cost, more desirable performance may be obtained. Difficulty arises in specifying what is a more desirable result: reduced overage or reduced shortage.

Decreasing the bin size may potentially improve the performance of the optimal policies, as the model is not forced to choose batched meal order adjustments. Depending on the distribution of differences, alternative values of alpha may also result in performance improvements.

Late penalty costs are small in comparison to the cost of a meal order. Thus, these costs effect the timing of the meal ordering where costs are not time dependent, rather than the final meal ordering performance. For example, the meal requirement may be ordered at decision epochs 3, 4, and 5 with equal costs. The late penalty cost causes the initial meal ordering to occur at decision epoch 4 rather than decision epoch 3.

The value of the terminal cost has the single greatest effect on the optimal policy and resulting performance. As shown in Figure 20, increasing the terminal cost shifts the distribution of provisioning error towards increased overage and reduced shortage.



Figure 20. Distribution of provisioning error over increasing terminal cost

We are interested in meal ordering policies that simultaneously minimize overage and shortage. However, such policies are only attainable when the change in passenger demand between decision epochs can be reliably predicted. Variability in model provisioning error is a direct result of variability in the distribution of passenger demand.

Historical passenger demand data exhibits considerable variability. Thus, model provisioning error is typically close to actual provisioning error in variability. Consequently, only marginal gains may be identified through use of the optimal policies. This illustrates two points:

- Given the variability in passenger demand, current meal ordering performance is close to optimal. There exists potential to achieve marginal gains through use of the optimal policies.
- The model may be evaluated over several scenarios to identify the relationship of overage versus shortage. Based on these results we may determine the expected cost of achieving a service level.

B. Overage versus Shortage

We seek to identify the relationship between overage and shortage by evaluating the model over a range of terminal costs. A policy based on a high terminal cost will result in reduced shortage but high overage, whereas, a policy based on low terminal cost will result in reduced overage and increased shortage. One may specify criteria for a maximum acceptable level of shortage and seek the policy providing the least amount of overage. The criteria may be of the form:

- The minimum cost policy that allows shortage up to five meals to occur on 1% of the flights.
- The minimum cost policy that never allows any shortage to occur on any flights.

Alternatively, one may specify a maximum acceptable level of overage, and seek the policy that minimizes shortage. The criteria may be of the form:

- The policy that allows overage up to five meals to occur on any flights, with the least shortage.
- The policy that allows overage up to ten meals to occur on 1% of all flights, with the least shortage.

The relationship between overage and shortage may be shown by plotting average overage versus proportion of flights experiencing shortage up to η meals. We use a specific measure for shortage (as opposed to average shortage) because the costs for each shortage are interpreted as large in comparison to the price of the meal. We evaluate the measures for overage and shortage over a range of terminal cost values. For a single flight, we obtain a group of points, which represent the curved relationship. Different curves are obtained for each value of η .

The relationship is described in the following example where $\eta=0$ (Figure 21).



Figure 21. Example average overage versus shortage curve

The average actual overage observed in the sample period was approximately 9.4 meals per flight. The actual proportion of short-catered flights was approximately 6.5%. By applying the policy obtained by the model, where average overage was close to 9.4 meals per flight, the proportion of short-catered flights would have decreased to approximately 3.3%. Alternatively, by applying the policy obtained where the proportion of short-catered flights was close to 6.5%, the average overage could have been reduced to

approximately 8 meals per flight. Here we see that the average overage associated with shortage occurring on 0% of the flights is 12.4 meals per flight, approximately 3 meals per flight higher than the actual average overage.

In this example, it was possible to simultaneously reduce the proportion of flights short-catered, while reducing the average overage. However, we can see that there is a limit to improved performance. According to these results no such policy exists to jointly reduce the average overage to 8 meals per flight and the proportion of flights short-catered to 1%.

In some cases the model does not generate meal ordering policies that outperform actual practice. This may occur as a result of one of the following:

- The modeling of the passenger demand was poor.
- The meal ordering clerk exercises exceptionally good judgement.
- The meal ordering clerk has access to information unavailable to the model. For example, the meal ordering clerk may receive a phone call regarding a flight cancellation on competitor flight.
- The historical pre-departure data is not representative, or a longer history is required.

Thus we may observe one of three possible scenarios in the results. These are presented in Figure 22 as a plot of average overage versus proportion of flights short-catered. Under the first scenario, actual practice outperforms the optimal cost ordering policies. Under the second scenario, actual practice closely matches performance obtained with the optimal cost ordering policies. The decision maker may choose to trade-off average overage to decrease the proportion of flights short-catered. Under the third scenario, the optimal cost ordering policies outperform actual practice. Here, the decision maker may choose one of three possible options:

- A policy that holds average overage constant, and decreases proportion of flights short-catered.
- A policy which holds proportion of flights short-catered constant, and decreases average overage.
- A policy that simultaneously decreases average overage and proportion of flights short-catered.



Figure 22. Possible model results

Where the meal ordering performance of the model is close to the performance achieved historically, the terminal cost value associated with the policy is an estimate of the 'in practice' terminal cost value. Note that this exercise is similar to performing a sensitivity analysis on the terminal costs.

C. Comparison of Current Performance to Model Performance

A group of 40 flights with a sufficient history was selected for modeling. These flights all travel from a single high activity center to one of 15 different destination stations. They are grouped into one of three groups depending on flight duration. Each flight was modeled separately, but no grouping was made for day of the week or season. While aircraft capacity may vary by flight route, it is constant across seasons and day of the week. The model parameters applied are consistent with those described in the previous section on *Basic Results* unless otherwise specified. Each flight was modeled over a series of terminal penalties spanning from \$5 to \$15,000 per short meal. The following bin sizes were applied based on aircraft capacity:

Capacity	Bin Size	Number of Flights
88	2	30
108	2.	6
180	3	2
236	4	2

Table 4. Bin sizes used in sample

In Figure 23 we present overage versus shortage curves pertaining to four flights at the level of zero shortage.





In the overage versus shortage curves shown in Figure 23, the square points represent the values obtained by evaluating the model at each specific terminal cost. Each circle represents the performance observed in actual practice for the particular flight. We assume that a policy exists between any two successive points. In 2 of the 4 flights, the model outperforms the historical performance, meaning that an optimal policy exists resulting in an overage and shortage pair more favorable than that observed in practice. Note that we seek a policy resulting in a pair closest to the origin. In the remaining 2 cases, the optimal policies result in meal ordering performance that closely matches actual practice.

The performance was evaluated for all 40 flights, at the service level where no shortage occurs. The results are presented in Table 5. In 21 cases out of the 40 flights, the optimal policies outperformed actual practice. In 8 cases, the optimal policies closely matched the actual practice, and in the remaining 11 cases, actual practice outperformed the optimal policies.

	Optimal policies outperforms actual practice		Optimal policies closely match actual practice		Actual practice outperforms optimal policies		
	N	Percent	N	Percent	Ν	Percent	Total
LONG	12	85.7%	1	7.1%	1	7.1%	14
MEDIUM	5	50.0%	5	50.0%	0	0.0%	10
SHORT	4	25.0%	2	12.5%	10	62.5%	16
	21	52.5%	8	20.0%	11	27.5%	40

Table 5. Model results by flight group

The performance obtained using the optimal policies is poor for short duration flights. In 10 out of 16 flights, the actual practice outperformed the optimal policies. Alternatively, the performance obtained with the optimal policies was good for long duration flights. The optimal policies outperformed actual practice in 12 out of 14 of these flights.

To determine the impact of applying the optimal policies we observe the point where the model equals or slightly outperforms the actual practice in proportion of flights with shortage. We multiply the model and actual average overage figures by a flight specific average cost per meal. Understandably, meal costs are generally higher for long duration flights. The results are provided in Table 6.

	Total Monthly Overage Cost								
N		Actual	Model	Difference					
LONG	14	29,466	24,228	5,238					
MEDIUM	10	6,842	6,762	80					
SHORT	16	5,811	6,625	-814					
	40	42,119	37,615	4,504					

Table 6. Impact of the optimal policies on overage costs

In the long duration flight group, the optimal policies result in overage costs almost 18% less than those observed in actual practice. Note that 3 of the 14 flights account for 73% of the cost improvement. In the short duration group, the optimal policies result in overage costs approximately 14% larger than those observed in actual practice. As shown in Figure 24 and Table 7, performance obtained with the optimal policies is also better in terms of proportion of flights short-catered.



Figure 24. Distribution of proportion of flights short-catered

	Number of F	Percent		
	Actual	Model	Difference	Difference
LONG	39	31	-8	-20.5%
MEDIUM	29	13	-16	-54.4%
SHORT	75	37	-37	-50.0%
	142	82	-61	-42.8%

Table 7. Impact of the optimal policies on number of flights short-catered

Recall, for each flight we choose the optimal policy that equals or outperforms actual practice in proportion of flights short-catered. Applying the chosen policies results in a decrease of over 42% in the number of flights short-catered. The majority of the improvement occurs in the medium and short duration flight groups.

Next, we seek the cost associated with achieving a pre-defined level of service. Two service levels are arbitrarily chosen: i) the average overage associated with no shortage, and ii) the average overage associated with shortage not exceeding 5 meals.

The actual average overage, current proportion of flights short-catered, and additional overage associated with zero shortage are presented by flight number in Table 8. Additional average overage is the model average overage minus the actual average overage. Negative values refer to savings in overage, either due to i) the optimal policies outperforming actual practice, or ii) the chosen service level is below the current service level. Similar figures are presented for the service level where shortage is less than 5 meals. Short, medium, and long duration flights have flight numbers beginning with 1, 2, and 3 respectively.

Actual proportion of flights

Additional average overage

	Actual			Where shortage	
Flight	average	Short catered by at		does not exceed 5	Where shortage
number	overage	least 5 meals	Short catered	meals	does not occur
1.01	8.2	1.7%	8.6%	0.9	4.2
1.02	5.4	1.7%	6.9%	2.1	5.3
1.03	5.1	1.6%	8.2%	-1.0	1.8
1.04	4.9	0.0%	10.5%	-1.2	3.3
1.05	4.7	5.3%	24.6%	4.4	6.7
1.06	4.6	6.6%	16.4%	3.0	5.1
1.07	4.6	0.0%	11.3%	1.8	2.5
1.08	4.2	4.1%	16.3%	0.5	4.1
1.09	4.1	6.8%	22.0%	10.3	15.5
1.10	4.0	0.0%	14.8%	-2.0	4.0
1.11	3.2	6.5%	19.4%	2.7	4.5
1.12	3.2	6.7%	18.3%	3.4	5.3
1.13	3.1	1.7%	6.8%	1.5	2.4
1.14	3.0	11.9%	30.5%	5.9	9.9
1.15	2.7	12.5%	35.7%	3.1	5.9
1.16	2.2	1.7%	10.0%	1.3	1.6
2.01	9.1	0.0%	6.3%	-3.1	-1.3
2.02	7.5	0.0%	1.9%	-1.6	0.1
2.03	7.2	0.0%	3.4%	-2.1	1.7
2.04	5.9	3.2%	9.7%	-1.4	3.7
2.05	5.5	3.3%	11.7%	1.2	2.6
2.06	4.9	1.7%	10.2%	-0.2	0.8
2.07	4.7	5.9%	19.6%	-0.4	3.4
2.08	4.5	2.2%	4.4%	-0.5	0.5
2.09	4.3	3.4%	15.3%	-1.3	2.0
2.10	2.8	0.0%	18.3%	-0.1	1.7
3.01	19.3	1.8%	5.4%	2.7	7.1
3.02	18.6	2.1%	4.2%	-2.2	3.7
3.03	14.4	2.2%	10.9%	2.8	4.9
3.04	12.6	15.6%	17.8%	2.0	9.4
3.05	10.4	0.0%	3.3%	-1.8	2.2
3.06	10.0	3.3%	10.0%	-0.9	1.5
3.07	9.3	3.3%	6.6%	-4.0	-1.3
3.08	8.1	3.3%	13.1%	1.6	6.8
3.09	7.6	0.0%	1.6%	-2.3	-0.6
3.10	7.2	8.3%	18.3%	2.4	7.9
3.11	7.0	0.0%	8.1%	3.2	4.8
3.12	6.7	6.7%	21.7%	1.2	5.5
3.13	6.4	0.0%	6.0%	3.9	4.8
3.14	6.2	0.0%	11.1%	-2.7	1.6

Table 8. Costs associated with achieving a zero shortage service level

In actual practice, many flights are short-catered, and often the shortage exceeds 5 meals. The short duration flight group exhibits the most shortage, and the long duration flight group exhibits the most overage. The additional average overage required to achieve a service level of no shortage is relatively high in comparison to actual average overage, exceeding 300% in some cases. As shown in Figure 25, as the actual proportion of short-catered flights increases, the additional average overage also increases.



Figure 25. Additional average overage versus actual proportion of flights short-catered

We observe the relationship between additional average overage and actual average overage. One would expect that a high actual average overage would require a smaller additional average overage to achieve a pre-defined service level. As shown in Figure 26, this relationship is not apparent across all flight groups.



Figure 26. Additional average overage versus actual average overage where shortage does not occur

Note that the short flight duration group has higher additional average overage, relative to their distribution of actual average overage.

Failure to provide a meal service on a long duration flight has a larger impact on customer goodwill, in comparison to a short duration flight. In practice the caterer tends to over-cater on long duration flights more so than on short duration flights. We can partially see this in Table 8 and Figure 26, where the short duration flights generally exhibit less actual average overage than the long duration flights. In Figure 27 we show the distribution of actual provisioning error for short, medium and long duration flights.



Figure 27. Distribution of actual provisioning error

The mean provisioning error for the short, medium and long duration groups are 3.3, 5.3, and 9.6 respectively.

Note that short and medium duration flight groups have a larger proportion of even-catered flights, in comparison to the long duration flight group. This may partially be a result of distribution of passenger demand for these groups. The short and medium duration flight groups have a large proportion of flights with final passenger load matching the aircraft capacity. Censoring of variability due to aircraft capacity facilitates the prediction of final meal order requirement.

Station-Wide Impact

For the 40 flights analyzed, the actual average overage is estimated to cost \$42,100 per month. The additional cost associated with achieving a level of service where no flight has a shortage exceeding 5 meals is estimated as \$2,500 per month. The total additional cost associated with achieving a level of service where no flight experiences a shortage is estimated as \$20,100 per month. The upper bound on this figure is obtained by calculating the additional cost of fully catering these flights, given that the actual overage still occurred. This figure is \$121,900 per month. These cost figures are based on actual average meal costs by flight number.

Estimating the system-wide impact of applying the optimal cost policies for all flights requires that the model be run for all flights over a sufficient range of terminal costs. In addition, to model the effect of seasonality, optimal policies must be separately tested by season. Thus, we require at least two years of pre-departure data, where the first year would be the model dataset, and the second year would be the test dataset. This would require a significant amount of data and processing time.

A crude estimate is provided based on meal costs. The group of 40 flights represents 59% of the total meals provisioned in a single high volume station. These flights represent only 46% of the total meal costs for that station. Assuming that the additional costs for the remaining 54% of total meal costs is linearly proportional to the additional overage costs estimated, we obtain a scale up factor of 2.16. The monthly station-wide additional cost estimates are provided in Table 9.

	40 Flights	Station-Wide
Estimated actual overage costs	42,100	90,900
Estimated additional overage costs - shortage less than 5 meals	2,500	5,400
Estimated additional overage costs - zero shortage	20,100	43,300
Additional overage costs - fully catered	121,900	263,100

Table 9. Monthly estimates of additional costs

Thus, we can see that achieving either of the two service levels presented requires that additional costs be invested into the system. Depending on the desired level of service, this cost may be substantial.

D. Sensitivity of Performance to Probability Density Function Specifications

The probability density function applied to the transition probability matrix is dependent on two factors: the bin size and the value of alpha. Recall, where the bin size is greater than one, we group together cells of the transition probability matrix accordingly. While this creates no loss of information in the matrices that are generated with a normal approximation of the differences, some information may be lost in matrices modeled with an empirical distribution.

The value of alpha is the relative weighting assigned to the empirical distribution, over the row-dependent normal distribution. A high alpha value relates to a high weighting on the empirical distribution, whereas a low alpha value relates to a high weighting on the normal distribution. This parameter represents the degree to which the empirical distribution is smoothed.

We observe the effects of different values of bin size and alpha for two flights. Flights 1.06 and 3.07 are chosen as they represent opposites in optimal policy performance in the results obtained previously. The actual practice outperformed the optimal policies in flight 1.06, whereas the optimal policies outperformed actual practice in flight 3.07. Both flights operate with the 88 seat capacity aircraft.

Optimal policies were obtained for the flights separately, over varying levels of terminal cost. The performance was evaluated in terms of average overage and proportion of flights short-catered. The values of bin size 1, 2, and 8 were applied over the values of alpha 0.1, 0.5, and 0.9.

The nine scenarios were compared using the following methodology. The distance from the origin to the actual proportion of flights short-catered, actual average overage coordinate was determined. In Figure 26, this distance is labeled z_a , and the direction vector v_a . Next, the distance from the origin to the intersection with the model overage versus shortage curve, along the vector v_a is obtained. This distance is labeled z_i .



Proportion of flights short-catered



A positive difference between z_a and z_i indicates that an optimal policy exists that outperforms actual practice. A negative difference indicates that actual practice outperforms the optimal policies. We compare the percent differences over the range of alpha and bin size values. The percent difference is obtained by dividing the difference by the value of z_i . Note that the proportion of flights short-catered is scaled up by 100 so that the average overage will not over-influence the results.

The results are presented in Tables 10 and 11 for flights 1.06 and 3.07 respectively. The actual practice outperforms the optimal polices for flight 1.06 over all combinations of alpha and bin size tested. We are interested in the cells with the lowest negative values of percent difference. The optimal policies for this flight perform best with the value of 0.9 for alpha. Also, policies where bin size is 8 appear to outperform the policies with lower bin size. This finding is unexpected as a large bin size results in a smoothed distribution, whereas a high alpha value results in high weighting on the empirical distribution.

Percent difference

			Bin Size			
		Alpha	1	2	8	
		0.1	-24.3%	-31.3%	-4.3%	7
		0.5	-12.7%	-18.9%	-4.6%	
		0.9	-4.1%	-5.3%	-4.1%	
1 1	10		• • •	C 1 1	11	

Table 10. Effects of varying levels of alpha and bin size: Flight 1.06

In flight 3.07, the optimal polices outperform actual practice over six combinations of alpha and bin size. The optimal policies closely match actual practice where alpha is 0.1 and bin size is 2 or 8. Where alpha is 0.9 and bin size is 8, actual practice outperforms the optimal policy. We seek the cells with the highest value of percent difference. The policies where alpha is 0.9 and bin size is 1 or 2 outperform actual practice by the largest amount. This may indicate that the empirical distribution holds some useful characteristics that we should not smooth out.

Percent difference

	Bin Size		
Alpha	1	2	8
0.1	3.7%	0.0%	0.0%
0.5	15.7%	21.0%	2.1%
0.9	26.2%	30.2%	-4.5%

Table 11. Effects of varying levels of alpha and bin size: Flight 3.07

These results indicate that optimal bin size and alpha values may be flight dependent, thus the analysis should be conducted over all 40 flights. A heuristic search for optimal bin size and alpha values may be added to a future version of the model. Generally, small values of bin size and high values of alpha appear to work well in these two dissimilar cases.

VI. CONCLUSION

In this thesis we have developed a finite horizon Markov decision process model for identifying minimum cost meal ordering policies. With this model, we investigated the multi-objective problem of overage and shortage, and obtained the costs associated with obtaining specific levels of service. Analysis was conducted on a sample group of 40 flights departing from a single high volume station.

According to the results obtained, current practice in meal ordering at Canadian Airlines often achieves performance close to that achieved with the optimal policies. Analysis by flight duration revealed that the optimal policies performed well on long duration flights and poorly on short duration flights. Medium duration flights exhibited performance close to actual practice. Overall, the model outperforms current practice on 52.5% of the flights, and closely matches the performance of actual practice in 20% of the flights. Applying the optimal policies at a level of service comparable to current practice yields cost savings of 17%, 1%, and -14% in the long, medium and short duration groups respectively.

Evaluating the model over varying levels of terminal cost generates overage versus shortage curves. Management may use these results to quantify the cost of achieving a service level, given the current processes. Depending on the level of service required by Canadian Airlines, the optimal policies may result in reduced costs.

Two levels of service were evaluated in terms of additional costs. The levels of service are i) where no flight experiences a shortage exceeding 5 meals, and ii) where no flight experiences any shortage. The station-wide monthly additional cost associated with achieving a level of service where no flight has a shortage exceeding 5 meals is estimated as \$5,400 per month. The station-wide additional cost associated with achieving a level of service where no flight experiences a shortage is estimated as \$43,300 per month.

The station-wide estimates are based on the scaled up results of the sample group. The accuracy of the estimates will improve by extending the analysis to all flights for the station. Similarly, accuracy may improve by evaluating the model over different seasons. A multivariate statistical analysis of the flight groups should be conducted to determine the appropriate resolution. Both improvements require more data and a significant amount of processing time.

Sensitivity analysis on the probability density function indicated that improvement in model performance might be achieved with different distribution parameters. Also, we observed that model parameters might be flight specific. A similar sensitivity analysis should be conducted over all flights to verify this hypothesis. The additional cost figures should be reassessed for the parameters that provide the best policy performance.

Given the potential to improve service level, we recommend that the additional analysis described above be conducted. Contingent on the model performance, Canadian Airlines may choose to implement the policies at an expected cost and benefit. Alternatively, if the choice of service level results in extremely high additional costs, Canadian Airlines may choose to invest the estimated additional costs towards process and logistics improvements.

Bibliography

Nahmias, Steven (1982). Perishable Inventory Theory: A Review, Operations Research, vol. 30, 680-708.

Puterman, M. L. (1992). *Dynamic Programming*, Encyclopedia of Physical Science and Technology, vol. 5, Academic Press.

Puterman, M. L. (1994). Markov Decision Processes: Discrete Stochastic Dynamic Programming, Wiley, New York.

Smith, Barry C., Leimkuhler, John F., Darrow, Ross M. (1992). Yield Management at American Airlines. Interfaces, vol. 22, 8-13.

Sun, Xiaoyun, Brauner, Erik, Hormby, Sharon (1998). A Large-Scale Neural Network for Airline Forecasting in Revenue Management, Kluwer Academic Publishers, 46-67, Dordrecht.

Appendix A. Overage versus Shortage Curves

We observe the overage versus shortage curves obtained for different values of η . Recall, η represents the maximum short catering that can occur on a single flight, measured in meals. As the value of η increases, the curves move closer to the origin. The circles represent the actual overage shortage points observed at the three levels of η in the historical data. In Figure A-1, we present curves obtained for flight number 3.08.



Figure A-1. Overage versus shortage with varying levels of η

In this example, the optimal policies closely match actual practice in performance, at the level where shortage does not occur. Where $\eta=2$ we see that the optimal policies clearly outperform actual practice. At the level of $\eta=5$ the optimal policies slightly outperform actual practice. These curves provide insight into the nature of the multi-objective overage versus shortage problem.

Appendix B. Effect of Passenger Demand on Subsequent Flights to the Same Destination

When two flights travel between a common origin-destination pair in a given day, one may expect information from the earlier flight to help predict changes in passenger load on the subsequent flight. For example, if the earlier flight was extremely overbooked the passenger load on the later flight may increase by the amount of passengers overbooked. Similarly, some passengers may have cancelled their tickets on the earlier flight, and booked a seat on the later flight. The effects of these scenarios are shown in the following example where flight B departs approximately 2 decision epochs after flight A. We denote the change in passenger load between intervals i and j as d_i.



Figure A-2. Scenario 1: passenger load shifts to the later flight

In this example, we see that the passenger load on flight A decreased by d_1 between decision epoch 1 and departure. Many of these passengers shifted to flight B resulting in a large increase in passenger load between decision epochs 3 and 2, denoted as d_3 .

Alternatively, in some cases passengers may attempt to travel on an earlier flight. This scenario is shown in the following example, where flight B departs approximately 2 decision epochs after flight A.



Figure A-3. Scenario 2: passenger load shifts to the earlier flight

In this case, the passenger load on flight A increased by d_1 between decision epochs 1 and departure. Most of these passengers were originally booked on flight B, thus the passenger load decreased by d_3 , which is close in magnitude to d_1 from flight A.

If such relationships exist between subsequent flights, we may potentially improve our prediction of change in passenger load for the second flight. From our two examples, we would expect d_1 and d_3 have a negative relationship. Depending on the time between departure for the two flights, the relationship may exist between other differences (i.e. between d_1 from flight A and d_2 from flight B).

A group of 40 flights were selected for analysis in the *Results and Sensitivity Analysis* section. Of these 40 flights, 25 had a flight preceding it by several hours, traveling to the same destination. We observe linear correlation between d_1 from the earlier flight, and the corresponding difference on the subsequent flight. We obtain values and significance levels based on flights between February 1998 and January 1999 inclusive.

The results are presented in Table A-1. Flight numbers beginning with 1, 2, and 3 refer to short, medium, and long duration flights respectively. Note the flights are intentionally numbered differently here to maintain the anonymity of the original data. In these results *lead time* refers to the difference in departure time between the subsequent flights. We are primarily interested in the differences in passenger load that correspond to this lead time. These are indicated as shaded cells in Table A-1.

		departure and:					Significance Level			
		d_1	d ₂	d ₃	d_4	d_1	d ₂	d ₃	d_4	
Flight Number	Lead time **	1 hour to departure	2 hours to 1 hour	3 hours to 2 hours	6 hours to 3 hours	1 hour to departure	2 hours to 1 hour	3 hours to 2 hours	6 hours to 3 hours	Alpha <= .05
1-A	2	0.15	-0.01	0.01	-0.05	0.00	0.82	0.83	0.32	N
1-B	2	0.01	0.05	-0.16	0.07	0.90	0.53	0.03	0.32	N
1-C	1	0.05	-0.08	0.01	0.00	0.34	0.12	0.79	0.94	N
1 - D	3	0.09	-0.06	0.02	-0.05	0.08	0.28	0.64	0.36	N
1 - E	2	0.23	-0.19	-0.20	-0.01	0.00	0.00	0.00	0.86	Y
1-F	2	0.07	-0.21	-0.04	-0.13	0.27	0.00	0.47	0.03	Y
1 - G	2	0.04	-0.08	0.05	0.05	0.51	0.15	0.37	0.36	N
1 - H	2	0.26	-0.18	-0.01	-0.02	0.00	0.00	0.84	0.69	Y
1-I	1	0.03	-0.08	0.05	0.00	0.59	0.11	0.38	0.97	N
1-J	4	0.02	-0.02	0.00	0.00	0.74	0.64	0.94	0.95	N
1-K	1	0.00	-0.14	-0.06	-0.10	0.99	0.01	0.27	0.05	Y
1-L	1	0.16	-0.17	-0.03	-0.01	0.00	0.00	0.59	0.82	Y
1-M	3	0.31	-0.03	-0.17	-0.07	0.00	0.52	0.00	0.20	Y
2-A	6	0.09	0.08	0.01	-0.02	0.13	0.15	0.85	0.69	N
2-В	4	0.09	-0.05	-0.09	-0.10	0.10	0.43	0.11	0.08	N
2-C	5	0.09	-0.04	0.03	-0.06	0.16	0.49	0.60	0.29	N
2-D	1	0.09	-0.09	0.04	-0.08	0.07	0.07	0.39	0.14	N
2-E	7	0.02	-0.13	0.10	-0.02	0.70	0.01	0.05	0.68	N
2-F	8	0.22	0.00	0.00	0.06	0.00	0.94	0.93	0.26	N
3-A	1	-0.09	-0.09	-0.05	0.00	0.18	0.19	0.51	0.99	N
3-В	1	0.06	-0.13	0.19	0.09	0.57	0.22	0.06	0.37	N
3-C	2	0.22	-0.06	-0.05	-0.03	0.00	0.29	0.33	0.53	N
3-D	2	0.09	-0.16	-0.09	-0.04	0.08	0.00	0.10	0.50	Y
3 - E	2	0.19	0.08	0.01	-0.03	0.00	0.15	0.81	0.52	N
3-F	6	0.23	0.01	0.12	0.09	0.00	0.82	0.04	0.14	N

Correlation between d₁ from previous

** Time since previous departure to same destination (in hours)

Table A-1. Correlation of passenger load differences between subsequent flights

We can see that the correlation vales are mostly small, ranging between -0.31 and 0.21. In only 7 of the 25 flights was the correlation significant. This finding forms the basis for assumption 1 presented in the *Model Formulation* section.

Appendix C. Modeling Differences in Passenger Load Between Pre-Departure Time Intervals with Seasonality

In modeling the differences in passenger load between pre-departure time intervals, we expect that incorporating seasonality may improve the accuracy of our estimates. We analyze two aspects of seasonality: i) day of week seasonality, and ii) standard seasonality (i.e. winter, spring, summer, fall). Recall, that we model the differences in passenger load between the final two decision epochs partially using a linear regression model, with *passenger load at 1 hour pre-departure* as the independent variable. The regression model is of the form:

$\begin{array}{l} \textit{Difference}_{1 \ \textit{hour pre-departure}} = \beta_0 + \beta_1 p l_1 + \varepsilon \\ \textit{to post departure} \end{array}$

Indicator variables are constructed to represent the days of the week and the seasons. These are incorporated in a multiple regression model with passenger load at 1 hour pre-departure. A stepwise variable selection technique is applied at a significance level of 0.15. This technique includes variables in the regression model to maximize the reduction in residual sum of squares, but re-evaluates the significance level of all included variables at each stage, removing variables with sufficiently large p values.

We separately analyze day of the week and standard seasonal effect, and then attempt a model including both indicators. Each flight is analyzed separately, and outlier observations are identified through studentized residuals obtained from a preliminary regression. Studentized residuals outside of the range of -3 to 3 are considered to be outliers, and the regression is rerun on all included observations.

The analysis was conducted on a sample group of 40 flights, previously selected for analysis in the *Results and Sensitivity Analysis* section. The flights are numbered 1, 2, and 3 representing short, medium, and long duration flights respectively.

Where the regression model includes day of week seasonality, the indicator variables range from Sunday to Friday inclusive. An indicator for Saturday is not included, as it would result in a multicollinearity problem. The day of week coefficients represent the difference in mean response between a given day of week and Saturday. For example, the coefficient β_3 measures the effect of Mondays on the one-hour to post departure differences in comparison to Saturdays. The regression model is of the form:

 $\begin{array}{l} Difference_{1 \ hour \ pre-departure} = \beta_{0} + \beta_{1}pl_{1} + \beta_{2}Sunday + \beta_{3}Monday + \dots \\ to \ post \ departure \\ \dots + \beta_{7}Friday + \varepsilon \end{array}$

In Table A-2 we present the regression coefficients obtained for each flight, where day of the week indicators were analyzed. We expect to see similar patterns among flights within the same group, however the results indicate that this is not the case.

	Passenger load							_ 2	R ²
Flight	at 1 hour pre-							R^2	Improve-
Number	departure	Sun	Mon	Tue	Wed	Thu	Fri	Original	ment
1.01	-0.171	-1.93	1.90		-8.89			0.171	0.256
1.02	-0.109	•			-	-	•	0.144	0.000
1.03	-0.140		1.90	2.88				0.263	0.037
1.04	-0.073	•		-1.01		•	-0.81	0.079	0.010
1.05	-0.115							0.192	0.000
1.06	-0.173		2.01			•	•	0.195	0.011
1.07	-0.205		•		104-23		1.18	0.263	0.003
1.08	-0.128	1.49		-1.18			-1.09	0.295	0.013
1.09	-0.162	3.31					1.67	0.168	0.027
1.10	-0.120	÷	0.97			•	•	0.366	0.004
1.11	-0.341		1	2.65	1.77	1.45		0.454	0.012
1.12	-0.228						1.40	0.342	0.003
1.13	-0.223	1.09			•	•		0.317	0.003
1.14	-0.167					3.09	2.12	0.178	0.037
1.15	-0.186				-1.43		1.86	0.258	0.016
1.16	-0.331		a desta de la	1.58	0.99	2.51		0.521	0.026
2.01	-0.203	-1.34	•	-1.93		1.44		0.121	0.027
2.02	-0.149		•	-1.44			2.27	0.351	0.024
2.03	-0.138		的主义的	1. 4. 14		1.49		0.146	0.009
2.04	-0.133							0.234	0.000
2.05	-0.297	•		-2.42	•	•		0.209	0.013
2.06	-0.123			-1.65		1.27		0.080	0.012
2.07	-0.117		1.09					0.280	0.003
2.08	-0.141			-1.48	1.91	1.35		0.085	0.036
2.09	-0.125			•			•	0.179	0.000
2.10	-0.073			-0.77				0.213	0.004
3.01	-0.137	-6.31	-5.16	-5.44	-5.45	•	-2.43	0.220	0.041
3.02	-0.064	-3.25			2.68			0.167	0.031
3.03	-0.154	3.84	2.82	•	•		•	0.224	0.014
3.04	-0.071			-4.50				0.080	0.022
3.05	-0.090	40.57			10.1	1.98		0.107	0.008
3.06	-0.152	1.81	-4.36	•			-4.21	0.246	0.070
3.07	-0.180							0.263	0.000
3.08	-0.182			-2.62				0.308	0.012
3.09	-0.203	1.59	1.1.1	74.11			2.10	0.321	0.013
3.10	-0.159	2.08			•	19.27		0.214	0.010
3.11	-0.120		-1.79		-1.22		-2.92	0.285	0.046
3.12	-0.118				-1.73	•		0.197	0.007
3.13	-0.194		-3.64	-4.02	•			0.170	0.033
3.14	-0.100			-0.96	-1.68	0.94		0.335	0.020
Counts	40	11	10	16	10	9	12	-	

Table A-2. Describing seasonality through regression coefficients

In these results R^2 original refers to the proportion of variability that was already explained by the *passenger load at one hour to departure* simple regression model. R^2 improvement refers to the additional variability explained with the additional day of week indicator variables. In both cases we use adjusted R^2 , which reduces the R^2 value to account for the number of regression parameters. This adjustment provides for measuring the reduction in the mean square due to the regression, thus the measure will not necessarily increase with the addition of superfluous variables.

The distribution of R^2 improvement is shown in the following boxplot. The values are less than 0.027 for 75% of the flights.



Figure A-4. Distribution of R^2 improvement from day of week indicators

We would also expect to identify a coarse day of the week effect, such as weekend or business travel days. As shown in the results, such an effect is not apparent.

Where the regression model includes standard seasons, the indicator variables include winter, spring, and summer. Again, an indicator for fall is not included, as it would result in a multicollinearity problem. The season coefficients represent the difference in mean response between a given season and fall. For example, the coefficient β_2 measures the effect of winter on the one-hour to post departure differences in comparison to fall. The regression model is of the form:

 $Difference_{1 hour pre-departure} = \beta_0 + \beta_1 p l_1 + \beta_2 Winter + \beta_3 Spring + \beta_7 Summer + \varepsilon_{10} to post departure$

In Table A-3 we present the seasonal regression coefficients obtained for each flight.

Flight	Passenger load				\mathbf{P}^2	R ²
Number	departure	Winter	Spring	Summer	Original	ment
1.01	0 166	w miter	2 163	2 025	0.171	0.044
1.01	-0.100	1.607	2.105	2.923	0.171	0.044
1.02	-0.117	-1.007	Discribert Bi		0.144	0.013
1.03	-0.130				0.205	0.000
1.04	-0.071	-1.465	•	•	0.079	0.000
1.05	-0.110	-1.405	1317		0.192	0.011
1.00	-0.102	3 575	3 /21	1 08/	0.195	0.007
1.07	-0.212	-5.575	-3.421	2 077	0.205	0.033
1.00	-0.121			2.077	0.295	0.019
1.09	-0.138	•	. 0.842	1 176	0.100	0.000
1.10	-0.111		0.842	1.170	0.300	-0.004
1.11	-0.338	-1.303	·		0.454	0.000
1.12	-0.220	1. A.			0.342	0.000
1.15	-0.210	•	0.948	-0.943	0.317	0.011
1.14	-0.160	Katura (1944)		2016/00/00/00/00/00/00/00/00/00/00/00/00/00	0.178	0.000
1.15	-0.175		-0.931		0.258	0.002
1.10	-0.346	-1.383	0.46 °		0.521	0.013
2.01	-0.197	×		-0.991	0.121	0.004
2.02	-0.134		1.409	• • • • • • • • • • • • • • • • • • • •	0.351	0.007
2.03	-0.151	-1.65/			0.146	0.009
2.04	-0.122			-1.209	0.234	0.009
2.05	-0.287	-1.6//		×	0.209	0.007
2.06	-0.127	•	-2.839		0.080	0.043
2.07	-0.121		-1.644		0.280	0.015
2.08	-0.158		18 - 18 - 18 - 18 - 18 - 18 - 18 - 18 -	-1./48	0.085	0.011
2.09	-0.123	1.205		·	0.179	0.012
2.10	-0.070	-0.762	-1.435		0.213	0.027
3.01	-0.126	-2.109	•		0.220	0.003
3.02	-0.057	-4.694			0.167	0.080
3.03	-0.178	-5.921	-3.823		0.224	0.042
3.04	-0.069	Marina da Marina		3.030	0.080	0.011
3.05	-0.092			2.206	0.107	0.017
3.06	-0.178	-4.16/		San Angelan	0.246	0.046
3.07	-0.180	×	-1.974	×	0.263	0.016
3.08	-0.172	A Statistica and a statistical statistical statistics and a statistical statistical statistics and a statistic			0.308	0.000
3.09	-0.200				0.321	0.000
3.10	-0.138			-1.397	0.214	0.007
3.11	-0.128	-1.139			0.285	0.007
3.12	-0.115	Chirdeala a - d	-1.309	• • • • • • • • • • • • • • • • • • • •	0.197	0.004
3.13	-0.174		•		0.170	0.000
3.14	-0.096			1. 新教育学校	0.335	0.000

Table A-3. Describing seasonality through regression coefficients

We expect to see flights within the same group to have coefficients on the same indicator variables. Again the coefficients do not follow an obvious pattern. The seasonal indicator variables do not drastically improve the amount of variability explained by the model. The distribution of R^2 improvement is shown in the following boxplot. The values are less than 0.015 for 75% of the flights.



Figure A-5. Distribution of R^2 improvement from seasonal indicators

Note that in one observation (flight 1.10), the R^2 improvement is a negative figure. This is primarily due to outlier removal, where an outlier observation included in the original model was removed from the model with seasonal variables.

A regression model with both day of week and seasonal indicators was also attempted. The following boxplot of R^2 improvement illustrates that the model produced results comparable to the regression model with day of week indicator variables.



Figure A-6. Distribution of R^2 improvement from day of week and seasonal indicators

Based on this methodology, there is no large improvement in modeling *difference in passenger load* between 1 hour pre-departure and departure by including day of week and seasonal indicators. This forms the basis for assumption 2 presented in the *Model Formulation* section.

Appendix D. Markov Decision Process Implementation

In this section we describe the development of the Markov decision process model and discuss how the optimal cost policies are obtained. Recall that we use backward induction to efficiently determine the minimum expected cost policy. We will describe the implementation starting with a one period problem, and expand the description to the remaining decision epochs. Recall that the decision epochs are numbered in descending order representing the stages remaining until the terminal stage.

Cost parameters are specified at the start of the model. Cost per meal, the time dependent return penalty, and a van delivery charge are determined based on historical cost data. A time dependent late penalty and a terminal cost are specified prior to running the program.

Based on the cost parameters we evaluate the terminal cost over all possible states. Thus we have a $\{pl_0\}\times\{mq_0\}$ matrix of dimension $\{Capacity+1\}\times\{Capacity+1\}$, which describes the cost associated with occupying a specific state in the at the point of departure. This matrix is referred to as V₀, where the (i, j)th entry represents the terminal cost given a flight departed with i passengers, and j meals. For example, the (80, 90)th entry would contain the costs associated with an overage of 10 meals.

Next, we evaluate the options that the decision-maker is faced with in the one-period problem. At decision epoch 1, the decision maker may adjust the meal order quantity from mq₁ (the current meal quantity) to mq₀, incurring an immediate cost. Between the last decision point, and the point of departure, the passenger load will change from pl₁ to pl₀ with probability $Pr(pl_0 | pl_1)$. Thus, at the terminal stage, the system will potentially occupy a state (pl_0, mq_0) incurring a terminal cost. We seek to choose the action that minimizes the immediate cost associated with adjusting the meal order quantity from mq₁ to mq₀, and the expected cost associated with the terminal stage.

The immediate costs of adjusting the meal order from mq_1 to mq_0 are calculated with the cost parameters and stored in a matrix of dimension $\{Capacity + 1\} \times \{Capacity + 1\}$ called CF₁. The (j, k)th entry represents the cost of adjusting the meal order from j meals in the current time interval, to k meals in the following time interval. The matrix represents the space of $\{mq_1\} \times \{mq_0\}$.

The probabilities of the passenger load changing from pl₁ in decision epoch 1, to pl₀ in the terminal stage are stored in a transition probability matrix labeled P₁. This matrix represents the space of $\{pl_1\} \times \{pl_0\}$ and thus has dimension $\{Capacity + 1\} \times \{Capacity + 1\}$. We obtain the expected cost associated with the terminal stage using both the P₁ and V₀ matrices. Multiplying P₁ by V₀ yields the expected subsequent cost matrix NC₁. The (i, k)th entry represents the expected cost of having a meal order of k meals in the terminal stage, when there are i passengers observed at decision epoch 1.



Figure A-7. Dimensionality of subsequent costs matrix

Using the immediate cost matrix CF_1 and the expected subsequent cost matrix NC_1 , we may combine the costs to determine the action for each state that minimizes the total cost. By adding combinations of rows

from each matrix, we obtain a 3 dimensional array. The $(i, j, k)^{th}$ entry represents the cost associated with occupying a state of i passengers and a meal order quantity of j meals in decision epoch 1, and adjusting the meal order so that there are k meals at the terminal stage. This matrix has dimension $\{Capacity + 1\} \times \{Capacity + 1\} \times \{Capacity + 1\}$ and contains all possible immediate cost and subsequent cost combinations.



Figure A-8. Dimensionality of all possible immediate and subsequent cost combinations

Having obtained all costs, we now seek the action (the meal quantity in the terminal stage mq₀) that minimizes total costs. Thus, for each state (pl₁, mq₁), we search all values of mq₀ for the minimum cost. The mq₀ value associated with this minimum represents the minimum cost action for a particular state. We store the minimum costs in a $\{pl_1\}\times\{mq_1\}$ matrix called MCC₁, and the actions mq₀ - mq₁ in a similar matrix called MCA₁. The matrix MCA₁ is the decision rule for decision epoch 1, describing the minimum cost action to take, given a current meal order quantity and passenger load.

Recall the value function as described in the Model Formulation section:

$$t = 5, ..., 1$$

$$v_{t}^{*}((mq_{t}, pl_{t}), mq_{t-1}) = \min_{a \in A} \left\{ r_{t}((mq_{t}, pl_{t}), mq_{t-1}) + \sum_{pl_{t-1} \in [0, Capacity]} \Pr(pl_{t-1} | pl_{t}) v_{t-1}^{*}(mq_{t-1}, pl_{t-1}) \right\}$$

$$t = 0$$

$$v_{0}^{*}((mq_{t}, pl_{t}), mq_{t-1}) = r_{0}((mq_{t}, pl_{t}), mq_{t-1})$$

$$a_{t}^{*} \in \arg\min_{a \in A} \left\{ r_{t}((mq_{t}, pl_{t}), mq_{t-1}) + \sum_{pl_{t-1} \in [0, Capacity]} \Pr(pl_{t-1} | pl_{t}) v_{t-1}^{*}(mq_{t-1}, pl_{t-1}) \right\}$$

In implementation, the V₀ matrix represents the values of v_0^* for all possible states. The immediate cost matrix CF₁ represents the immediate cost/reward function $r_t((mq_t, pl_t), mq_{t-1})$ for all possible actions. The subsequent cost matrix NC₁ represents the expected cost function $\sum_{t=0}^{t} \Pr(pl_{t-1} | pl_t) v_{t-1}^*(mq_{t-1}, pl_{t-1})$ evaluated over all possible states. The function of combining $pl_{t-1} \in [0, Capacity]$

immediate costs and expect subsequent costs in the matrix MCD_1 , represents the contents within the main brackets of the value function (before minimization). The minimum cost matrix MCC_1 represents the

resulting value function v_1^* for decision epoch 1, evaluated over all possible states. The minimum cost action matrix MCA₁ represents the associated minimum cost actions a_1^* , for all possible states.

To extend this over the remaining decision horizon we generate the necessary cost and transition probability matrices, and recursively evaluate the value function from decision epoch 2 through 5. Thus for decision epoch 2, we construct the immediate cost matrix CF_2 based on the cost parameters. We then multiply the transition probability matrix P_2 by the minimum cost matrix (or value function) for decision epoch 1, MCC₁, generating the subsequent cost matrix NC₂. By combining the immediate cost matrix CF_2 and the expected subsequent cost matrix NC₂ we obtain MCD₂. Choosing the minimum action in MCD₂ will generate the minimum cost matrix MCC₂. We continue to construct and evaluate the value functions, MCC_n until we reach decision epoch 5.

Our optimal policy π will consist of decision rules d₅, d₄, d₃, d₂, and d₁ which are stored in matrices MCA₅, MCA₄, MCA₃, MCA₂, and MCA₁ respectively.