Coffee, Tea, or ...?: A Markov Decision Process Model for Airline Meal Provisioning

Jason H. Goto
Centre for Operations Excellence
University of British Columbia
209 - 2053 Main Mall,Vancouver, BC Canada V6T 1Z2
jgoto@coe.ubc.ca, 604-822-1800

Mark E. Lewis
Department of Industrial and Operations Engineering
University of Michigan
1205 Beal Avenue, Ann Arbor, MI 48109-2117
melewis@engin.umich.edu, 734-763-0519

Martin L. Puterman
Faculty of Commerce and Business Administration
University of British Columbia
2053 Main Mall, Vancouver, BC Canada V6T 1Z2
marty@coe.ubc.ca, 604-822-8388

February 13, 2002

Abstract

This paper develops and analyzes a finite horizon Markov decision process model for the airline meal provisioning activity focusing explicitly on developing policies for determining and revising the number of meals to upload. Using one year of daily data from over 40 flights, the paper shows that the optimal policies can result in both improved customer service and significant dollar savings, especially in long haul flights. It also applies the model to derive an efficient frontier and investigate tradeoffs between having too few and too many meals on a flight.

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 (Smith, Leimkuhler, and Darrow, 1992) and optimization of crew schedules (Vance, Barnhart, Johnson, and Nemhauser, 1997). The majority of these applications focus on core airline operations. A nice general reference for operations research in this industry is Yu (1998). 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, significant costs, and has direct impact on customer service.

The inflight meal provisioning process involves producing meals in an airport kitchen by a caterer and delivering them to a plane for eventual inflight passenger service. A fundamental question is how to determine the number of meals to be prepared so as to ensure a high level of passenger service while keeping costs as low as possible. We assume the following decision-making scenario. Well in advance of departure, the caterer determines the number of meals to produce for a specific flight. At several subsequent decision points prior to the departure of a flight, the caterer may adjust the meal quantity to be delivered to the plane. Delivery and production costs vary over time with unscheduled deliveries close to departure more costly and constrained by van capacity. Due to inherent variability, the final passenger load is not known with certainty prior to departure. In fact, in the data we have analyzed the final number of passengers on a flight varies from those booked in at one hour prior to departure by as much as $\pm 10\%$.

The purpose of this paper is to develop a decision-making framework for this process based on Markov decision processes (MDP) and illustrate its potential impact by applying it to historical data. Our representation of the meal provisioning operation is based on processes at Canadian Airlines where this study was carried out. We believe that these processes are typical of meal provisioning at a wide range of carriers and that the results herein have wide applicability. More details are available in Goto (1999).

The problem analyzed herein has its roots in stochastic inventory control. It may be viewed as a newsvendor model (c.f. Gallego and Moon (1993)) with demand forecast updates and multiple quantity adjustment opportunities. In essence, the decision maker chooses a meal order quantity based on early demand information (the number of booked passengers) and as updated information becomes available (the number of booked passengers changes) adjusts the order quantity. As is common in fashion good inventory models, late adjustments to the order quantity are more costly. This occurs in our problem because late meals must be delivered by a special van. The final cost structure is similar to the newsvendor model, there are costs for overage and underage but the good has zero salvage value. The most closely related paper to ours is that of Eppen and Iyer (1997) which considers the optimal management of a fashion inventory system using early order information to adjust the order quantity. Closely related is a fashion good supply chain application described by Fisher and Raman (1996). From an alternative perspective, our model provides an approach for updating forecasts of the final order quantity based on accumulated data. Such forecast models originate with Hausman (1969) and have subsequently been considered by Graves (1986), Heath and Jackson (1994), Toktay and Wein (2001) and Aviv (2001). The first three papers provide a dynamic model for forecast updating and the latter two papers characterize optimal policies. Our work differs from these in that we model the demand evolution as a Markov chain and use dynamic programming to take this information into account. Also, our primary focus is application so that we are primarily concerned with computing optimal policies and not deriving structural results.

The remainder of the paper is organized as follows: Section 1 describes the model. We formulate the model as a Markov decision process in Section 2. Our approach to applying the model to the problem at hand is discussed in Section 3. The results of our study and the conclusions are contained in Sections 4 and 5, respectively.

1 Problem Description

Meal provisioning activities may be divided into two main stages; a production stage and an adjustment stage. Production primarily takes place in the flight kitchen and involves the preparation, assembly and refrigeration of meals for a specific flight. At a subsequent point of time the meals are removed from refrigeration and delivered in a catering truck to the flight. Adjustment involves altering the meal quantity after it has left the kitchen. Process details follow.

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 volume requirements 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 estimates the meal quantity requirement for the flight based on the number of tickets booked for the flight, a forecast of the passenger load based on the flight history, a forecast of the required meal quantity (everyone on board may not want a meal or some may require a special meal), and his/her own personal experience.

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. 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 and it is 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 left the kitchen area. Adjustments are made with a limited capacity van, which travels between the flight kitchen and the aircraft. The driver verifies the adjustment with in-flight staff and carries the meals to the aircraft. An order may require a late adjustment when the passenger load changes substantially after delivery of the initial meal order. These late adjustments are costly and inconvenience passengers since the driver often has to carry large meal carriers up the aisles of a crowded plane.

For the flights considered we assume the key decision points for the production process occur at eighteen, six, and three hours prior to departure. In practice decisions to alter the production quantity may be made at intermediate time intervals. Subsequent decisions are made throughout an adjustment period in the remaining time prior to departure. See Figure 1.

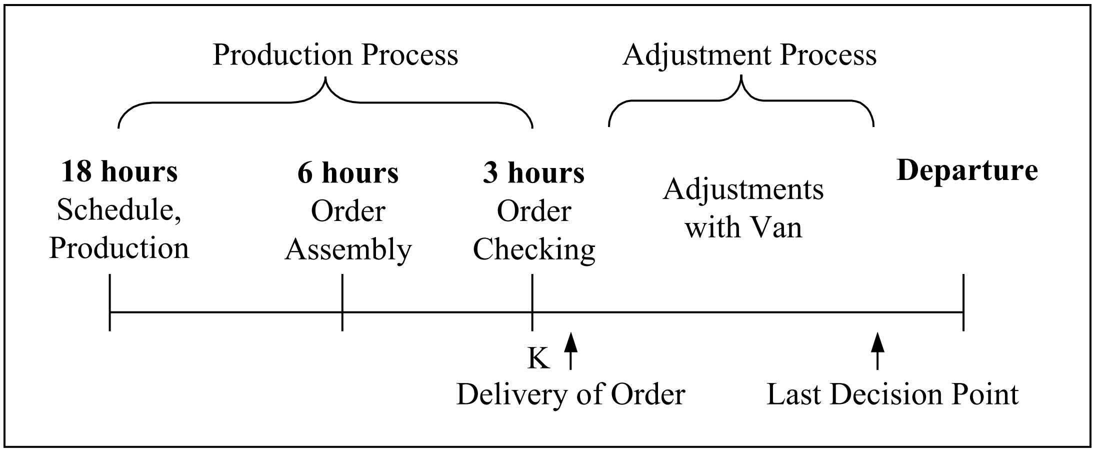

Figure 1: Timeline of events in the provisioning decision process.

We assume production, adjustment, shortage and overage costs. Production costs include all food and labor that are directly involved in preparing and assembling a meal order. They vary directly with meal quantity. An order adjustment that occurs prior to delivery incurs no penalty cost, subsequent adjustments made with a limited capacity van incur a per meal penalty cost plus a fixed delivery charge. Costs of overage are taken to be the cost of meals in excess of passenger load. Shortage costs 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. This cost is difficult to quantify but it is clearly higher on a long haul flight than on a short haul flight. We also measure the level of customer service as the fraction of passengers who do not have a meal available (even if they do not wish to consume it). In our analyses below we use customer service level as a performance metric and not a constraint.

2 Markov Decision Process Model Formulation

As noted above, the meal ordering problem has some features of a newsvendor problem, but given the multiple decision points and changing information, we formulate it as a finite horizon Markov decision problem (c.f. Puterman (1994)). *For simplicity we assume a single meal type and a single passenger class.* The expression *flight* refers to an aircraft allocated to a specific route at a specific time with a unique flight number. Let

$$q_t = \text{Meal quantity at decision epoch } t$$

$$l_t = \text{Booked and standby passenger load at decision epoch } t, t = 1, \ldots, N$$

$$l_0 = \text{Actual passenger load at departure}$$

$$M = \text{Aircraft capacity in passengers}$$

$$K = \text{Last decision epoch prior to meal delivery}$$

$$N = \text{Total number of decision epochs.}$$

A Markov decision process is characterized by a state space, an action set, decision epochs, costs, and transition probabilities. For our model they are as follows:

- The state space: $S = \{0, 1, \ldots, M\} \times \{0, 1, \ldots, M\}$. The first element of the state space represents the *meal quantity*; the current number of meals that have been allocated to a given flight. Depending on the decision epoch, this may be in production, in a holding refrigerator, or on-board the aircraft. The second element is the *passenger load*; the number of booked and stand-by passengers. At the point of departure this is the number of on-board passengers. Note that we bound the first component of S by the aircraft capacity M. This quantity could in fact be increased to take overbooking into account. We do not explore that possibility here.

- Decision epochs: $T = \{0, 1, \ldots, N\}$ where time $t \in T$ represents the *remaining* time until the end of the planning horizon; the last time a meal can be delivered to an aircraft. Decision epochs need not be evenly spaced and represent critical decision points. No decision is made at decision epoch 0; it corresponds to the departure time and is included to evaluate the outcome of the meal allocation process.

- Action set: Denote the set of actions available in state s at decision epoch t by $A_{s,t}$. At each decision epoch the decision-maker chooses the number of meals to have available at the following decision epoch. During the production phase, when $t = N, \ldots, K$, the set of available actions is $A_{s,t} = \{0, 1, \ldots, M\}$. The action represents the order "up to " or "down to" quantity. For the adjustment phase, that is for $t = K - 1, \ldots, 1$, the set of available actions is $A_{\{q,l\},t} = \{max(q - V, 0), \ldots, min(q + V, M)\}$ where V denotes the van capacity. Note that during the adjustment phase, alterations to the allocated meal quantity are limited by van capacity.

- Costs: The costs $c_t(s, a)$ are defined for $s \in S$ and $a \in A_{s,t}$ for $t = 0$ by

$$c_0(q_0, l_0) = b(l_0 - q_0)^+ + e(q_0 - l_0)^+ \tag{1}$$

and for $t = 1, \ldots, N$ by

$$c_t((q_t, l_t), a) = variable\ meal\ cost_t + return\ penalty_t$$
$$+ van\ delivery\ charge_t + late\ penalty_t. \tag{2}$$

The terminal cost c_0 penalizes overages and underages; b denotes the per unit shortage cost and e the per unit overage cost. Variable meal costs refer to the cost per meal. If meals are returned after the order has been delivered to the aircraft a per meal penalty cost is incurred. For example, a meal that costs \$10 may be removed from the meal order without penalty when $t \geq K$. However, for $t < K$ a penalty of 50% of meal cost is incurred on any returned meal. Similarly, an increase in the meal order quantity during the adjustment stage results in a fixed van delivery charge. The late penalty represents the caterers' preference for building the meal order earlier than later and is included as an adjustment factor. This cost increases as the departure time draws near.

- Transition Probabilities: The transition probabilities give the likelihood of changes in pre-booking and stand-by levels between decision epochs or between pre-booking and stand-by level at the final decision epoch and load at departure. The only stochastic element in the transition probabilities is the change between l_t and l_{t-1} which is outside the control of decision maker. Changes in q_t depend deterministically on the action a. That is $q_{t-1} = a$. Thus $p_t((q_{t-1}, l_{t-1})|(q_t, l_t), a)$ depends only on

l_t, l_{t-1} and a. We denote the transition probability for the load by $P_t(l_{t-1}|l_t)$ so that

$$p_t((q_{t-1}, l_{t-1})|(q_t, l_t), a) = \begin{cases} P_t(l_{t-1}|l_t) & \text{if } q_{t-1} = a, \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

A (deterministic) decision rule, $d_t(q, l)$, specifies the number of meals to "order up to" or "reduce down to" at decision epoch t given that the meal quantity is q and the passenger load is l. A policy $\pi = \{d_N, d_{N-1}, \ldots, d_1\}$ is a sequence of decision rules and gives the decision-maker a prescription for making decisions at each decision epoch in each state. For a particular policy π, the expected total cost is

$$v_N^\pi(q_N, l_N) = \mathbb{E}^\pi_{(q_N, l_N)} \left\{ \sum_{t=1}^{N} c_t((q_t, l_t), d_t(q_t, l_t)) + c_0(q_0, l_0) \right\}, \tag{4}$$

where the expectation is conditional on the initial meal quantity q_N and passenger load l_N. Note that the only stochastic element in (4) is the passenger load. The *value function* of the MDP, v^* satisfies

$$v_N^*(q_N, l_N) = \inf_{\pi \in \Pi} v_N^\pi(q_N, l_N), \tag{5}$$

for all $(q_N, l_N) \in S$, where Π is the set of all (history dependent, randomized) policies. Since the state and action space are finite, we are guaranteed the existence of a Markovian, deterministic policy that achieves the infimum in (5) (cf. (Puterman, 1994, Proposition 4.4.3)).

The value of this MDP may computed by solving the finite horizon optimality equations

$$v_t(q_t, l_t) = \min_{a \in A_{(q_t, l_t), t}} \{c_t((q_t, l_t), a) + \sum_{l_{t-1} \in \{0, 1, \ldots, M\}} P_t(l_{t-1}|l_t) v_{t-1}(a, l_{t-1})\} \tag{6}$$

for $v_t^*(\cdot)$, where $t = 1, 2, \ldots, N$, $(q_t, l_t) \in S$, An optimal decision rule achieves the minimum in (6) for all states with v_t replaced by v_t^*. The standard solution approach is backward induction.

The model formulated above assumes that:

1. The passenger demand on one flight does not affect demand on subsequent flights. Transitions in passenger demand between the pre-departure decision epochs are modelled independently on a flight-by-flight basis, thus any information regarding other flights is not considered.

2. All meals ordered in a pre-departure time interval are ready and delivered in the following time interval with certainty.

3. The number of meals on each aircraft does not affect the demand for seats on the current flight.

3 Model Application

We now describe our approach to applying the MDP model. We choose 5 decision epochs to reflect practice. The first three decision epochs are 36 hours, 6 hours, and 3 hours pre-departure. The final two decision epochs are at 2 and 1 hour pre-departure. In our notation this corresponds to $N = 5$ and $K = 3$.

We apply it to 40 flights with the same origin. Flights are classified into short haul, medium haul and long haul based on destination. For each flight, we use one year of daily data. To gain insight to how this approach may perform in practice we use *cross-validation*. Pre-booking and final load data was divided into two contiguous subsets of 6 months each. We refer to the first subset as the *training data* and the second as the *test data*. The training data was used to parameterize the model and derive an optimal policy. The test data was used to assess its performance. We use the expression *provisioning error* to refer to the difference between the final meal quantity and passenger load; positive values correspond to excess and negative values to shortages.

The following performance measures were evaluated for the test data:

- The provisioning error distribution
- The percentage of flights that experience a shortage
- The percentage of flights that experience a shortage or overage exceeding 5 meals
- The average overage and shortage for flights on which there were overages or shortages, respectively

Other possible performance measures include the number of adjustments to the meal order or the number of van deliveries.

3.1 State Aggregation

The state space consists of all possible meal quantity and passenger load pairs. As capacity increases, the state space size increases quadratically and the number of actions increases linearly. For example, a 9 seat model (10 states including the zero state) requires 1000 action evaluations at each decision epoch, whereas a 19 seat model requires 8000 action evaluations. A 108 seat capacity aircraft (Airbus A320) with 5 decision epochs requires approximately 10 minutes of processing time on a Pentium 200 MHz PC. A typical fleet can have aircraft with as many as 380 economy class seats (Boeing 747). For this reason, we chose to analyze an aggregate model in which a seat in the model represents multiple seats in practice. We refer to the number of actual seats represented by a single seat in the model as the *bin size*. In the results section costs are scaled so that output can be compared to practice. Algorithms were coded and all calculations were carried out using the SAS software package. Complex calculations and matrix operations were performed using SAS/IML and data analysis and report preparation used SAS/STAT and SAS/GRAPH.

3.2 Transition Probability Estimation

As noted above, the transition probabilities are characterized by the conditional probability distribution of changes in passenger load between decision epochs, $P_t(l_{t-1}|l_t)$. We require estimates of these probabilities between five pairs of time points: 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. Clearly the number of quantities to be estimated is large, for example in a ten state aggregate model for a single flight we must estimate 90 probabilities (since the probabilities in each row must sum to 1) for each decision epoch or a total of 450 probabilities while a 108 seat aircraft would require 57780 entries. Our approach to estimating these probabilities combines a parametric model for load changes with a direct estimate based on observed transition frequency; (see Figure 2). This approach both fills in missing cells and smooths out the distribution across rows. The available data for estimating these probabilities are daily passenger loads at 5 pre-departure time points and the final passenger load for each flight on each day.

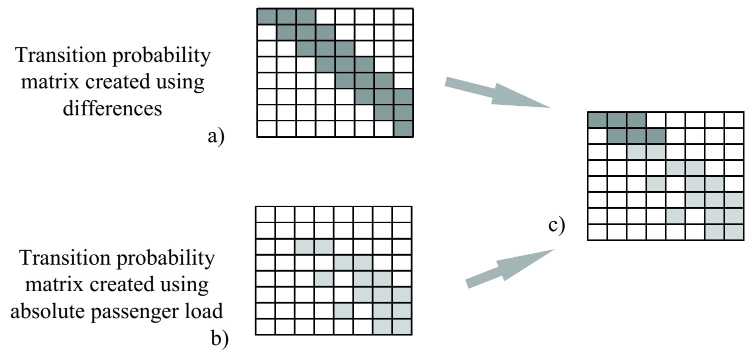

Figure 2: Generating a transition matrix by combining direct estimates with estimates of changes.

3.2.1 Direct Estimation of Transition Probabilities

At each decision epoch and for each initial load, we first estimate the probability of each subsequent load as its relative frequency. This is the maximum likelihood estimate of these probabilities. We refer to these estimates as *empirical transition probabilities*. The methodology captures passenger load dependent transitions, however, an extremely large number of observations is required to accurately estimate all transition probabilities. In practice, this approach produces an excessive number of zero entries, as transitions between many state pairs are never observed in our data.

3.2.2 Estimation of Changes in Passenger Loads

We now discuss estimation of transition probabilities based on changes in passenger load between decision epochs. Figure 3 shows the distributions of load changes between the six time points derived from data for a single-stage flight with a 108 seat capacity aircraft, observed over a six month period. Observe that the distributions of differences are centered near zero with the exception of the last transition which corresponds to load changes from one hour pre-departure to post departure. This distribution has a negative mean with considerable variability. This may be due to last-minute ticket cancellations, late arrivals or overbooking with some booked passengers being shifted to other flights. Another explanation is also offered below. *This extensive variability in the final period is the key factor that makes determining the appropriate meal quantity to upload difficult.*

Before developing a parametric model for load changes we investigate the relationship between the load

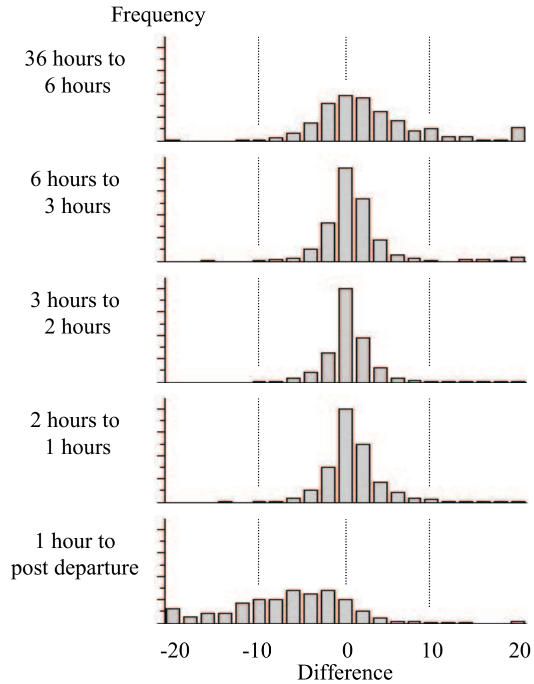

Figure 3: Distribution of differences between successive pre-departure intervals.

at a decision epoch and the change in load until the next decision epoch. The distribution of changes may become more variable as passenger load increases, or the mean of the differences may shift. In Figure 4, boxplots show the relationship between load changes between one hour pre-departure and post-departure and the passenger load at one hour pre-departure for one particular flight. The solid line represents a linear regression of load difference on passenger load. It shows a decreasing trend; the greater the load the more likely a decrease in load.

Figure 4 also shows that the differences are constrained by capacity. The line labelled "Capacity" represents an upper bound on load changes during the last period, so that the on-board passenger load does not exceed capacity. For example if the observed one hour pre-departure load is 100 in a 108 seat aircraft, then the changes in the last period are constrained to fall between $+8$ and -100 so it is apparent that the mean will be negative. As Figure 4 also shows, this effect is further exaggerated by overbooking. When the one hour pre-departure load exceeds capacity, which occurs for all loads in Figure 4 at which the line labelled "Capacity" is below 0, then the only allowable changes are negative.

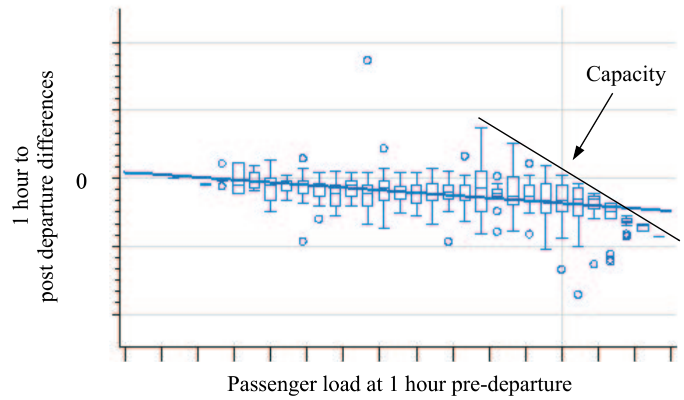

Figure 4: Boxplots of load changes as a function passenger loads at one hour pre-departure. Axes scales are unlabelled to respect data confidentiality agreement.

In most studies, ignoring censoring in the data leads to inefficient and biased estimation. Wu and Hamada (2000, Chapter 12.3) discuss several ways to compensate for this. Ding et. al (2000) show how to combine censoring with optimization in a dynamic inventory model with lost sales. For our model we take censoring into account by allowing the change distribution at the last decision epoch to depend on the load.

We note that a relationship between load and load change was only significant at the final decision epoch. For other decision epochs the two measures did not appear to be correlated. Consequently in deriving parametric models for load changes, only that at the last decision epoch accounted for the dependence on the load.

3.2.3 Normal Approximation

The distributions of the differences were fit with a shifted Poisson distribution and a normal distribution, independent of passenger load for each decision epoch except the last. Both distributions provided a comparable and reasonable fit. Based on this observation, we decided to use a normal distribution to model passenger load changes. We used different approaches for the first four decision epochs ($t = 2, 3, 4, 5$) and the final one ($t = 1$). In all cases we refer to these estimates as *normal transition probabilities*.

We first describe models for the first four decision epochs. To remove the effect of outliers we computed a trimmed mean and standard deviation based on deleting the upper 10% and lower 1% of the data. We chose to delete more observations from the upper tail because most extreme observations were observed to be positive (see Figure 3). These large positive values were attributable to late flight arrivals or cancellations.

Let m_t and s_t denote the trimmed mean and standard deviation respectively for each difference, where $t = 2, 3, 4, 5$. Recall that l_t is the load at time t and takes on discrete values. Furthermore, note that the lower time index refers to a decision epoch **closer** to departure. The normal transition probabilities were computed by

$$P_t(l_{t-1}|l_t) = \begin{cases} \Pr(Y_t \le -l_t) & l_{t-1} = 0, \\ \Pr(Y_t \le l_{t-1} - l_t) - \Pr(Y_t \le l_{t-1} - l_t - 1) & 0 < l_{t-1} < M, \\ \Pr(Y_t > l_{t-1} - l_t - 1) & l_{t-1} = M, \end{cases} \tag{7}$$

where the random variable Y_t represents the change in load between t and $t - 1$ and is assumed to follow a normal distribution with mean m_t and standard deviation s_t for $t = 2, 3, 4, 5$. By using trimmed estimates of the mean and standard deviation, the normal distribution represents the majority of the data well.

For the last decision epoch, $t = 1$, we allow the change distribution to depend on the load. We model the relationship between l_1 and l_0 as a simple linear regression of the form:

$$\text{Difference} = \beta_0 + \beta_1 l_1 + \epsilon \tag{8}$$

where β_0 and β_1 are regression parameters and ϵ is the error term. The normal transition probabilities are computed using (7), with 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 when the number of observations is large and the passenger load is within the range of the data used in the regression. A more precise estimate would use the prediction error standard deviation.

3.2.4 Combined Estimation

The two transition probability estimates were combined to fill in missing cells and smooth out distributions. When there were sufficiently many observations we used a weighted combination of the normal transition

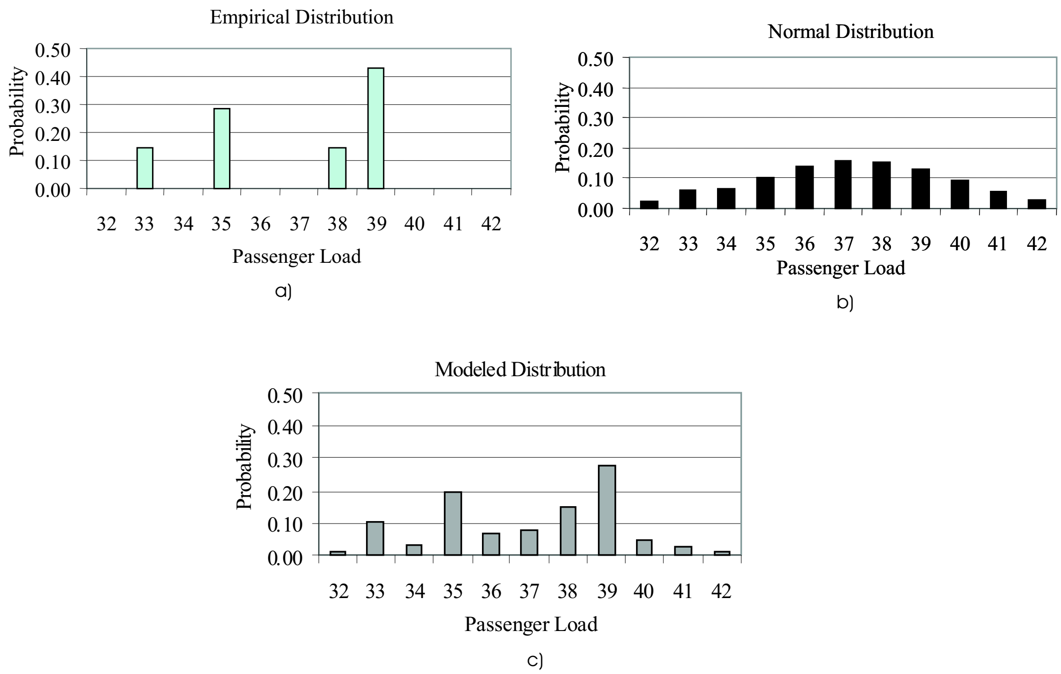

Figure 5: Filling empty cells in the transition matrix ($\alpha = .5$)

probability and the empirical probability density functions as follows.

$$\text{Estimated Probability} = \alpha \times \text{Empirical Transition Probability} + (1 - \alpha) \times \text{Normal Transition Probability} \tag{9}$$

where $0 \leq \alpha \leq 1$. For example, suppose a given row of the frequency table consists of seven observations starting with 39 passengers at time t and becoming $\{33, 35, 35, 38, 39, 39, 39\}$ at time $t - 1$. The empirical distribution, the estimated normal distribution and a pooled estimate of the two distributions with $\alpha = .5$ appear in Figure 5.

Rows with 5 or fewer observations are replaced with the corresponding estimates obtained using the differences as described in Section 3.2.2.

We investigated the effects of different values of bin size and α for two flights. The results indicate that the effect of bin size and α values on policy performance may be flight dependent. Generally, small values of bin size and high values of α appear to yield results with low average overage and a low proportion of flights short-catered.

3.2.5 Use of Explanatory Factors

An analysis was conducted to determine if any exogenous factors could help model the change in passenger load between the final two intervals. These transitions are the most critical in the performance of the model since there are higher costs associated with meal ordering during the adjustment process. 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 (i.e. flight 1 from A to B is cancelled so that flight 2 has a higher passenger load)
- 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 model and a predictive model was obtained using standard variable selection methods. 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 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.

Also, a wide range of time series methods were applied to determine whether the change in load between two decision epochs on one day could be predicted using data from previous days. This approach also did not appear to be useful.

4 Results

In this section we describe results of applying the model. We discuss the structure of the optimal policies, evaluate the tradeoff between overage and shortage, investigate the effect of flight duration, and estimate the cost of obtaining a desired service level. Recall that the expression *flight* refers to a specific route at a specific time. Usually this corresponds to a unique flight number. To apply the model to a particular flight requires estimating model parameters on the training set, determining an optimal policy and then applying it to the test data set at the 5 decision epochs each day. Results must then be summarized appropriately.

The model was applied using historical data for a selected 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 data set into a training data set and a test data set. An α value of 0.5 was used in (9) equally weighting the empirical and normal transition probabilities. The cost per meal varied between \$2 and \$12 depending on flight duration. The per meal shortage cost was chosen as \$120 and the overage charge was set equal to the cost per meal. The late penalty costs are \$0, \$0, \$2.50, \$2.50, and \$7.50 for the decision epochs 5 through 1, respectively and the van delivery charge was set equal to \$25. The cost of removing a meal from a plane during the adjustment period was set to half the meal value. The aircraft capacity was set at 108 seats, and a bin size of 9 was applied, so that the aggregate model had 12 seats. All references to the passenger load and meal ordering quantity pertain to the aggregate model unless otherwise stated. In practice, a much smaller bin size would be used.

4.1 The Optimal Policy

We describe some properties of the optimal ordering policy for a typical flight. At decision epoch 5, 36 hours pre-departure, the optimal action in all states is to order zero. This seems reasonable since there will be other opportunities to order meals at no cost before the adjustment process. At decision epoch 4 which represents 6 hours prior to departure the optimal policy is to order a meal quantity corresponding to the passenger load up to load 3; if the load is 4 order 5 meals; if the load is 5 to 8 order 7 meals and if the (pl, mq) epoch 3

	0	1	2	3	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	3	3	3	3	3	3	4	3	3	3	3	3	3
3	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

a)

(pl, mq) epoch 2

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

b)

(pl, mq) epoch 1

	0	1	2	3	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	3	2	3	4	5	6	7	8
2	3	3	3	3	3	3	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

c)

Figure 6: Optimal ordering policy for the last three decision epochs.

load exceeds 8 order 10 meals. The optimal policies in the last three decision epochs are more complex and represented in Figure 6. The decision rules are presented in 13 by 13 tables, where the row represents the current passenger load, the column represents the current meal order quantity and each cell contains the meal quantity after the adjustment (the action taken). Entries with the same value are shaded in the same grey-tone. (Note these displays are much more effective in color.)

Observe that at 3 hours pre-departure (decision epoch 3), the policy adjusts the meal order quantity to match the passenger load except in a few cases, for example when the passenger load is 8 and there are 5 meals on order, the order should be increased to 7. Note that time point 3 represents the last opportunity to adjust the meal order without incurring a van delivery charge, or a return penalty cost.

At 2 hours and 1 hour pre-departure the meal order is adjusted via van delivery so that a fixed delivery charge is incurred for every order in these two decision epochs. The van is assumed to have a fixed capacity of 4 meals, the effect of which is apparent in the decision rules at these two epochs. For example at decision epoch 2, if there are 10 meals on board and the passenger load is zero, then the number of meals on board is

	Provisioning Error	
Description	Actual	Model
Mean	9.81	7.99
Standard Deviation	8.46	6.96

Table 1: Distribution of provisioning error over 120 days for a typical flight: model versus actual.

reduced to 6. As the model chooses the set of decisions that minimize the expected reward over all decision epochs, if an order is to be placed at decision epoch 2, it may select a meal order adjustment that reduces the likelihood of having to incur two fixed delivery charges. Furthermore, note that at the final decision epoch when the passenger load is less than 5, the optimal policy orders one more meal than required as a result of trying to avoid the large shortage cost. However, when the load is higher, the negative slope of the regression causes the model to compensate for potential cancellations; no extra meals are ordered.

Optimal policies for other groups of flights generally follow the same form as this example with slight variations due to differences in the distribution of passenger demand. While records of the meal ordering quantities at each decision epoch were not readily available, interviews with catering staff revealed that the policies were similar to actual meal ordering practices.

4.2 Model Performance

We now apply decision rules obtained from the training data set to the four month test data set. Recall that we refer to the final meal order quantity minus the passenger load at departure as the *provisioning error* and compare it for the model and actual practice. We first consider a typical flight. Table 1 shows that both practice and the optimal policy over-cater the flight but the optimal policy results in a lower sample mean and variance over a 4 month period than was observed in practice.

We now compare the optimal policy to practice for the 40 selected flights. In Table 2 optimal policies derived from the model with parameters estimated in the training data set and the meal quantities used in practice are compared to actuals for each day in the test data set. Average overage denotes the total overage observed in the test data for flights on which there were excessive meals divided by the total number of flights with an excessive number of meals. As shown in Table 2, the average overage obtained with the

Description	Provisioning Error Actual	Model
Average overage (Meals)	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 2: Performance indicators for 40 selected flights: model versus actual.

optimal policies is lower than the average overage observed in practice (8.33 versus 10.19 meals). The optimal policies also produce a lower proportion of flights with overage and shortage exceeding 5 meals. This is particulary encouraging since the model ordering quantities are constrained to be multiples of the bin size. With the bin size set at 9, the only allowable meal ordering quantities are 0, 9, 18, ... while in practice any order quantity is possible. We see that the optimal policies result in meal ordering policies that outperform practice on several dimensions. We expect improved results with smaller bin sizes. We quantify this potential savings below.

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 α may also result in performance improvements. However, the considerable variability (see Figure 3) in the demand data limits the potential improvement.

4.3 Overage/Shortage Tradeoffs

In this section we discuss using the model to investigate the tradeoff between shortages and overage. In particular, we attempt to translate the shortage cost to a service measure.

The trade-off curve in Figure 7 depicts the efficient levels at which the system can operate. It was derived for a representative flight from the optimal policy obtained by varying the per meal shortage penalty cost from \$5 to \$15,000. This curve illustrates the trade-off between average overage and the proportion of flights at which there were fewer meals than passengers. It serves as a benchmark for assessing current practice as well as an illustration to management of possible cost-service trade-offs. It can be interpreted as follows. Suppose management policy requires that no more than 6% of flights be short-catered. Then

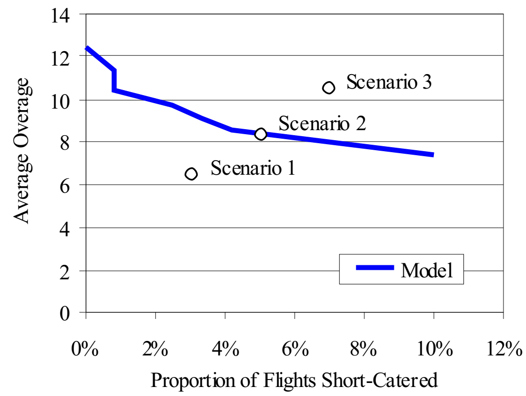

Figure 7: The efficient frontier for a typical flight.

this corresponds to a shortage cost at which the optimal policy produces an average overage of 8 meals per flight.

Also included on the graph are points indicating performance measures for three possible situations:

- Scenario 3 in which practice was inefficient and could be enhanced using an optimal policy.
- Scenario 2 in which practice replicated optimal performance.
- Scenario 1 in which practice outperformed the optimal performance as determined by the model.

If Scenario 3 occurs, the system is not performing optimally and can be improved in several ways. For example, the overage level of roughly 10 meals can be retained and the proportion of short-catered flights can be reduced to about 1%. Scenario 1 may appear puzzling at first since it indicates that actual practice outperforms the optimal policy. This occurs because in practice, the meal order clerk may have information that is not available to the model such as other flight cancellations or a large volume of unregistered standby passengers, or alternatively, the demand distribution or bin size may not be well chosen.

4.3.1 Effect of Flight Duration

We now investigate the relationship between flight duration and the frequency with which the three scenarios identified in the previous section occur.

A group of 40 flights of varying duration with sufficient data was selected for modelling. These flights departed from a single high activity center to one of 15 different destination stations. They were classified into three groups on the basis of flight duration. Aircraft capacity varied across flights but remained constant on each flight throughout the study period. Model parameters were set at values described above except where noted.

The chosen bin sizes appear in Table 3. An overage/shortage tradeoff curve was derived for each flight

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

Table 3: State aggregation in sample flights.

using the approach described above. The average real performance for that flight was also plotted on the tradeoff curve and we observed which of the three scenarios that flight fell into. Results are presented in Table 4. In 21 out of 40 cases the optimal policies outperformed practice. In 8 out of 40 cases, the optimal policies closely matched practice, and in the remaining 11 cases, practice outperformed the optimal policies. Observe that in long haul flights, the optimal policies outperformed current practice in 85% of the cases. Only in short haul flights were the model results dominated by practice. This suggests that there is a significant opportunity to apply the model to the more costly long haul flights.

We now estimate the economic value of the potential improvements suggested by the model. For each

	Optimal policies outperform practice		Optimal policies closely match practice		Practice outperforms optimal policies		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 4: Comparison of practice and model performance by flight duration. Entries are the number of flights and percentage of row totals.

		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 5: Impact of the optimal policies on overage costs.

flight we obtain the overage level on the efficient surface corresponding to the proportion of flights observed to have too few meals. The reduction in overage is another summary of the potential savings obtainable using model results. The model and actual average overage figures are multiplied by a flight specific average cost per meal to get the total monthly overage costs. Understandably, meal costs are generally higher for long duration flights. The results are provided in Table 5. In the long duration flight group, the optimal policies result in overage costs almost 18% lower than those observed in 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% greater than those observed in practice.

The optimal policies also performed better when the proportion of flights short-catered is considered. 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.

4.3.2 Costs of Obtaining Specified Service Levels

Next we seek to estimate the cost associated with achieving a pre-defined level of service. Two service levels are arbitrarily chosen; no shortage and shortage of at most 5 meals.

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 to be $2,500 per month. The total additional cost associated with achieving a level of service where no flight experiences a shortage is estimated to be $20,100 per month. The upper bound on this figure is obtained by calculating the additional cost of fully catering flights with shortages, given that the actual overage still occurred. That is to say, when a flight is short we immediately send a van to make up the

	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 6: Estimates of monthly costs to obtain specified service levels.

difference and when a flight has too many meals, the cost of those meals are a loss. This figure is $121,900 per month. These cost figures are based on actual average meal costs by flight number and are tabulated in Table 6.

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 data set, and the second year would be the test data set. This would require a significant amount of data and processing time. A rough 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 directly proportional to the additional overage costs estimated, we obtain a scale factor of 2.16. The monthly station-wide additional cost estimates are also provided in Table 6. Thus, we see that achieving either of the two service levels presented requires an additional investment. Depending on the desired level of service, this cost may be substantial.

5 Conclusions

In this paper we have developed a finite horizon Markov decision process for identifying minimum cost meal ordering and adjustment policies. We used the model to investigate the possibility of developing policies that simultaneously reduced overage and shortage, and obtained the costs associated with obtaining specific levels of service. Analysis was conducted on a sample group of 40 Canadian Airlines flights departing from a single high volume station.

According to the results obtained, using optimal policies could potentially reduce both provisioning error and its variability. Analysis by flight duration revealed that the optimal policies obtained using the Markov decision process model approach enhanced performance considerably on long duration flights and did not appear to be beneficial for short duration flights. Optimal policies for medium duration flights exhibited performance close to 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/shortage tradeoff 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 the airline, 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 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 multi-variate 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 assessed for the parameters that provide the best policy performance.

Results of this analysis were well received by Canadian Airlines management. They were especially impressed with the potential to operate meal provisioning more efficiently on long haul flights. In addition, they viewed the model as an effective approach for decision-making since it added structure to a rather subjective decision-making process. Unfortunately, implementation was delayed pending a restructuring of inflight operations as a result of the merger of Canadian Airlines and Air Canada.

6 Acknowledgements

This research was carried out in The Centre for Operations Excellence at UBC with support from the Natural Science and Engineering Research Council and the MITACS National Centres of Excellence program. We also wish to thank the staff of Canadian Airlines for access to information and facilities and cheerful cooperation.

References

Y. Aviv, "Effect of forecasting capabilities on supply chain coordination," (2001), working paper.

X. Ding, M. L. Puterman, and A. Bisi, "The censored newsvendor and the optimal acqusition of information," *to appear, Operations Research* (2002).

G. D. Eppen and A. V. Iyer, "Improved fashion buying with Bayesian updates," *Operations Research* 45, 805–819 (1997).

M. L. Fisher and A. Raman, "Reducing the cost of demand uncertainty through accurate response to early sales," *Operations Research* 44, 87–99 (1996).

J. H. Goto, *A Markov decision process model for airline meal provisioning*, Master's thesis, Faculty of Commerce, University of British Columbia, Vancouver, BC (1999).

S. C. Graves, "A tactical planning model for a job-shop," *Operations Research* 34, 522–533 (1986).

G. Gallego and I. Moon, "The distribution free newsboy problem; review and extensions," *The Journal of the Operational Research Society* 44, 825–834 (1993).

W. Haussman, "Sequential decision problems: a model to exploit existing forecasters," *Management Science* 16, B93–B111 (1969).

D. C. Heath and P. L. Jackson, "Modelling the evolution of demand forecasts with application to safety stock analysis in production/distribution systems," *IIE Transactions* 26, 17–30 (1994).

M. L. Puterman, *Markov decision processes: discrete stochastic dynamic programming*, John Wiley and Sons, Inc., New York (1994).

B. C. Smith, J. F. Leimkuhler, and R. M. Darrow, "Yield management at American Airlines," *Interfaces* 22, 8–13 (1992).

L. B. Toktay and L. M. Wein, "Analysis of a forecasting-production-inventory system with stationary demand," (2001), preprint.

P. H. Vance, C. Barnhart, E. L. Johnson, and G. L. Nemhauser, "Airline crew scheduling: a new formulation and decomposition algorithm," *Operations Research* 45, 188–200 (1997).

C. J. Wu and M. Hamada, *Experiments: planning, analysis, and parameter design optimization*, Wiley Series in Probability and Statistic, Wiley and Sons, Inc., New York (2000).

G. Yu, "Optimization applications in the airline industry," in *Handbook of combinatorial optimization*, 635–726, Kluwer Academic Publishers, Boston, MA (1998).

Figure 1: Timeline of events in the provisioning decision process.

Figure 2: Generating a transition matrix by combining direct estimates with estimates of changes.

Figure 3: Distribution of differences between successive pre-departure intervals.

Figure 4: Boxplots of load changes as a function passenger loads at one hour pre-departure. Axes scales are unlabelled to respect data confidentiality agreement.

Figure 5: Filling empty cells in the transition matrix ($\alpha = .5$)

Figure 6: Optimal ordering policy for the last three decision epochs.

Figure 7: The efficient frontier for a typical flight.

Table 1: Distribution of provisioning error over 120 days for a typical flight: model versus actual.

Table 2: Performance indicators for 40 selected flights: model versus actual.

Table 3: State aggregation in sample flights.

Table 4: Comparison of practice and model performance by flight duration. Entries are the number of flights and percentage of row totals.

Table 5: Impact of the optimal policies on overage costs.

Table 6: Estimates of monthly costs to obtain specified service levels.