Airline Overbooking Models with Misspecification

Kannapha Amaruchkul*· • Patipan Sae-Lim

School of Applied Statistics

National Institute of Development Administration (NIDA)

118 Serithai Road, Bangkapi, Bangkok 10240, Thailand

January 10, 2010

Abstract

Static overbooking models are studied. Suppose that each reservation shows up independently, and that the probability of showing up is identical among all reservations. Then, the random show demand follows the binomial distribution. However, in practice some overbooking modules assume that the show demand is the product of the overbooking level and the random show-up rate. The decision model embedded in a commercial revenue management system is misspecified. In this article, we explore the consequences of the modeling error. Through numerical experiments, we find that the performance of the model with misspecification decreases as the show-up probability decreases. Among our three choices of show-up rate distributions, the deterministic model performs worst, whereas the beta model performs best. The normal show-up model performs better than the deterministic model, but worse than the beta model.

Keywords: Revenue Management; Overbooking; Stochastic Model Applications; Operations Research

*Phone: 066-2-727-3070. Fax: 066-2-374-4061. E-mail: kamaruchkul@gmail.com

1 Introduction

Overbooking is practiced by nearly all passenger airlines. They may accept more reservations than their fixed capacity in order to compensate for cancellations and no-shows, which could be as high as 50% (Smith et al., 1992). The financial gain from the overbooking practice is one billion US dollars or more (Bailey, 2007). The overbooking module becomes indispensable to commercial revenue management (RM) systems.

Generally speaking, the objective of the overbooking model is to find an *overbooking level/limit*-the maximum number of reservations to hold at any time-that minimizes an expected total cost. The expected total cost is calculated with respect to the probability distribution of the *show demand* (*show-ups*), the total number of reservations that survive to the time of services. The total cost is comprised of an *oversale cost*, which occurs if the realized show demand exceeds the capacity, and a *spoilage cost*, which occurs if the realized show demand is less than the capacity.

The functional form of the show-ups can affect the overbooking level recommended by the model. Two models with different specifications of the show demand may not lead to the same expected total costs, and they may yield different overbooking levels. Models in practice commonly assume that the show demand is linear in the overbooking level; i.e., given the overbooking level x, the number of show demands is xR, where the random variable R is referred to as the *show-up rate*. Some commercial RM systems assume the normal/Gaussian show-up rate distribution, whose the mean and variance are periodically estimated from historical data (Popescu et al., 2006).

Although the above specification that the show demand equals the product of the overbooking level and the show-up rate is simple and quite prevalent in practice, it is theoretical incorrect under certain conditions. Suppose that (i) each reservation shows up independently, and that (ii) the probability of showing up is identical among all reservations. Then, the show demand given the overbooking level x follows a binomial distribution with parameters (x,θ), where θ is the show-up probability. Under conditions (i) and (ii), the linear assumption in the airline's decision model is incorrect; we say that a *model misspecification* occurs, and that the airline makes a *modeling error*.

In the RM practice, there is an iterative process in which the control (e.g., the overbooking level) from the optimization model is enacted, the data (e.g., the realized show demands) are collected over several flights, the parameters (the mean and variance of the show-up rate distribution) are forecasted, and finally the new control is determined from the optimization model given the updated parameters. In this article, we want to explore the consequences of the modeling error that the optimization model is misspecified.

In the optimization model, we consider three show-up rate distributions, namely normal, beta, and deterministic. For each of the three misspecified models, we provide a closed-form expression for the overbooking level. To benchmark and evaluate these models, we construct a model, in which the show demand theoretically follows a binomial distribution. We also obtain an optimal overbooking level with respect to the benchmark (binomial) model. To study the behavior of the iterative process with the misspecified optimization model, we

perform a series of numerical experiments. We find that as the iterative process goes on for a long time, the sequence of the average costs with the given misspecified model converges almost everywhere. The long run average cost is greater than the optimal expected cost with the binomial model. In all tested problem instances, the long-run average cost with the deterministic show-up rate is worse than that with the normal distribution, whereas the latter is worse than that with the beta distribution. All three overbooking models are not robust to the modeling error.

Overbooking models can be broadly categorized into two types, namely dynamic and static models. In the static models, the dynamics of reservation requests and customer cancellations over time are ignored. In the dynamic models, such inter-temporal effects are explicitly accounted for. An overview of the overbooking problem can be found in, e.g., Chapter 4 Talluri and van Ryzin (2004), and Chapter 9 Phillips (2005).

The dynamic model is often formulated as a Markov decision process. Examples of the dynamic overbooking problems are e.g., Chatwin (1992) and Subramanian et al. (1999). In Subramanian et al. (1999), the booking horizon is divided into a number of decision periods; in each period a booking request from a certain fare class may arrive, or a reservation may be cancelled, nothing happens. If the booking request arrives, then the decision is whether or not to accept the request. If the other two events happen, then no decision is needed. In the terminal period, the expected cost associated with no-shows is incurred. The distribution of the show demand is assumed to follow a binomial distribution. The objective is to maximize the expected total net contributions over the booking horizon and the terminal period. Bertsimas and Popescu (2003) study the network dynamic overbooking model, in which each itinerary may require more than one legs to get from an origin to a destination. There are other extensions to the dynamic overbooking problem, e.g., the inclusion of the multiple reservation classes in Karaesmen and van Ryzin (2004).

In this article, we do not study the dynamic overbooking problem and consider only the static overbooking model, since it is similar to the overbooking module in most commercial RM systems. As in Subramanian et al. (1999), the classical static overbooking model assumes that the show demand follows a binomial distribution. Thompson (1961) finds that the binomial distribution adequately fits the data collected from Tasman Empire Airways. Details of the binomial static overbooking model can be found in, e.g., Section 4.2.1 Talluri and van Ryzin (2004). Our benchmark model resembles the one in Talluri and van Ryzin (2004).

Unlike the binomial model, several static overbooking models assume that the show demand is the product of the overbooking level and the show-up rate. This approach is found in e.g., Kasilingam (1997), Popescu et al. (2006), and Luo et al. (2009). The random show-up rate can be modeled using a parametric distribution, such as uniform (e.g., Kasilingam 1997), beta (e.g., Luo et al. 2009), and normal. Popescu et al. (2006) argue that modeling the show-up rate as the normal random variable, which is quite common in practice, in not appropriate. They use a nonparametric method and obtain a histogram, in which the number and size of bins are constructed based on a wavelet method. In these articles, static overbooking problems are studied alone without the iterative process. In ours, the parameters of the show-up rate distribution are iteratively updated.

There is substantial literature in econometrics and statistics on model misspecification. For instance, Wu (1973) and Hausman (1978) propose

"specification tests" for detecting if the regression model is misspecified. In Presnell and Boos (2004) and Hansen et al. (2006), the decision maker hypothesizes that its model is misspecified, carries out the specification test, and generates the new model if there is insufficient evidence to reject the null hypothesis of misspecification. Unlike these papers, we do not endow the airline with the specification test. Such extension would be an interesting future research direction.

There are few operations research papers on model misspecification. Cooper et al. (2006) examine consequences of model misspecification in the two-class passenger RM problem. The airline makes incorrect assumptions about customer behaviors and chooses its optimal booking limit according to the Littlewood's rule. The paper shows that the modeling error leads to the so called spiral-down effect. The problem studied in Cooper et al. (2006) is the allocation problem not an overbooking problem, whereas ours is the overbooking problem

The rest of the paper is organized as follows. In Section 2, we present and analyze the overbooking models. The iterative process is described in Section 3. In Section 4, we report the results of numerical experiments, and we conclude with some thoughts on future research directions in Section 5. Proofs are in the Appendix.

2 Overbooking Problem

We consider a static overbooking problem. Because of their simplicity, such models become the basis of the most widely used methodology for making overbooking decisions (Talluri and van Ryzin, 2004). Define an overbooking problem as determining an overbooking level so that the expected total cost is minimized. Since the airline operates many repeat flights, we can assume that the decision maker is risk neutral, and the objective of minimizing the expectation is appropriate.

Throughout this article, $\mathbb{N}$ denotes the set of natural numbers, and $\mathbb{Z}$ the set of integers. For a real number y, $[y]^+ = \max(y, 0)$ the positive part of y, $\lfloor y \rfloor = \max\{n \in \mathbb{Z}; n \leq y\}$ the floor of y, $\lceil y \rceil = \min\{n \in \mathbb{Z}; n \geq y\}$ the ceiling of y. Assume that the capacity is a known constant c. If an overbooking level is set to x, denote the random show demand as $S(x)$. Let $a_o \geq 0$ be the per-unit oversale cost, and $a_s \geq 0$ the per-unit spoilage cost. The expected total cost is given as:

$$\widetilde{f}(x) = E\big[a_o[S(x) - c]^+ + a_s[c - S(x)]^+\big] \qquad (1)$$

In (1), the first and second terms are the oversale and spoilage costs, respectively. The oversale cost is computed as the per-unit oversale cost a_o times the number of show-ups that are denied boarding $[S(x) - c]^+$. The spoilage cost is found similarly. Consider the following problem:

$$\min \{f(x) = (a_o + a_s)E[(S(x) - c)^+] - a_s E[S(x)]\} \qquad (2)$$

Since $\widetilde{f}(x) = f(x) + a_s c$, an optimal overbooking level that minimizes $\widetilde{f}(x)$ is identical to the one that minimizes $f(x)$. Henceforth, Problem (2) is studied.

The airline chooses an overbooking level that minimizes the expected cost, which is calculated with respect to the distribution of the show demand $S(x)$. In practice, it is usually the case that the airline does not know the actual distribution of the show demand, but it makes overbooking decisions based on *perceived* models. We will shortly describe some perceived models, whose variants are embedded in some commercial RM systems. To evaluate and benchmark the perceived models, we also develop the *actual* model, in which the distribution of the show demand is known. We adopt similar terminology as in Evans and Honkapohja (2001) pp. 28-30.

2.1 Actual Model

Suppose that $x \in \mathbb{N}$ reservations have been made in advance and that each reservation requires a single seat (no group booking). Conditions (i) and (ii) are assumed to hold. With the actual model, the show demand $S_0(x)$ follows a binomial distribution with parameters x and θ. Let $f_0(x)$ be the objective function in (2), where we replace $S(x) = S_0(x)$.

Proposition 1. *With the actual model, the objective function $f_0(x)$ is convex on* N. *An optimal overbooking level is given as*

$$x_0^* = \text{argmax } \{x \in \mathbb{N} : (a_s + a_o)P(S_0(x-1) \geq c) \leq a_s\} \qquad (3)$$

The optimality condition in (3) can be explained intuitively as follows. Given that the current overbooking level is $x-1$, we want to know whether to overbook one more seat. We would incur an oversale cost, if the show demand from the current reservations [of $(x-1)$ seats] is at least the capacity. Hence, the expected marginal oversale cost is $a_o P(S_0(x-1) \geq c)$. We would incur a spoilage cost, if the show demand from the current reservations is strictly less than the capacity. Thus, the expected marginal spoilage cost is $a_s P(S_0(x-1) < c)$. If the expected marginal spoilage cost is at least the expected marginal oversale cost [i.e., $a_o P(S_0(x-1) \geq c) \leq a_s P(S_0(x-1) < c)$], then we would overbook one more seat.

2.2 Perceived Model

Journal articles on overbooking problems (e.g., Popescu et al. 2006, and Kasilingam 1996) suggest that airlines typically do not use a sophisticated approach to predict the show demand. It is commonly assumed that the show demand is linear in the overbooking level. Specifically, if the overbooking level is equal to x, then the show demand is xR, where R is the show-up rate. The distribution of the show-up rate is constructed from historical data.

We restrict our attention to parametric methods and consider three distributions that the airline might use to model the show-up rate. R_i be the random

show-up rate in perceived model i. R_1 has a degenerate distribution; i.e., $P(R_1 = \rho) = 1$, where $\rho \in (0, 1)$ represents a deterministic show-up rate. R_2 has a normal distribution with mean μ and variance σ^2, and R_3 follows a standard beta distribution with shape parameters a and b.

The deterministic model is documented in e.g., Talluri and van Ryzin (2004) p.147 and Phillips (2005) p. 213. A variant of the normal model is implemented in practice (Popescu et al., 2006). The beta distribution is used to model the show-up rate of air-cargo shipments (Luo et al., 2009). The support of the standard beta distribution is the open unit interval $(0, 1)$, whereas that of the normal distribution is the whole real line $(-\infty, \infty)$. Since the show-up rate should lie between zero and one, the beta model is more theoretically sound than the normal model.

Suppose that the airline uses perceived model i. Let y_i denote the parameter of the show-up distribution R_i; i.e., $y_1 = \rho$, $y_2 = (\mu, \sigma^2)$, $y_3 = (a, b)$. Let $f_i(. \mid y_i)$ be the objective function in (2), where $S(x) = x R_i$. Denote $\tilde{x}_i(y_i) = \arg\min f_i(x \mid y_i)$. Let $g_i(. \mid y_i)$ be the probability density function of R_i for each $i = 2, 3$.

Proposition 2. *For each i, the objective function $f_i(x \mid y_i)$ is convex in x. With perceived model 1, $\tilde{x}_1(y_1) = c/y_1$. With perceived model $i \in \{2,3\}$, $\tilde{x}_i(y_i)$ solves*

$$\int_{\alpha_i}^{c/\tilde{x}_i(y_i)} t g_i(t \mid y_i) dt = \frac{a_o}{(a_s + a_o)} E[R_i] \qquad (4)$$

where $\alpha_2 = -\infty$, and $\alpha_3 = 0$

Obviously, the left-hand side in (4) is decreasing in $\tilde{x}_i(y_i)$. The overbooking level can be found via a classical search procedure. The solution $\tilde{x}_i(y_i)$ found in the proposition may be not be an integer. If an integer overbooking level is desired, one could set it to $\tilde{x}_i(y_i) = \arg\min \left\{ f_i(\lfloor \tilde{x}_i(y_i) \rfloor \mid y_i), f_i(\lceil \tilde{x}_i(y_i) \rceil \mid y_i) \right\}$.

3 Iterative Process

Suppose that the airline operates many repeat flights. With each perceived model, an "optimal" overbooking level depends on the show-up rate distribution, whose parameters are periodically forecasted from the historical data. As new data become available, the airline updates the parameters of the show-up rate distribution, the overbooking level is chosen with respect to the updated distribution, and the process continues. These are sometimes referred to as the iterative data collection-forecasting-optimization process.

Suppose that the airline updates information every m flights. Define the t-th decision period to be a time in which the t-th forecast becomes available. Each period of the process consists of three steps: optimization, data collection, and

forecasting, respectively.

At the beginning the *t-th* period, the forecast for three perceived models are $y_{1t} = \rho_t$, $y_{2t} = (\mu_t, \sigma_t^2)$, $y_{3t} = (a_t, b_t)$. In the optimization step, the overbooking model i is $x_i^*(y_{it})$. In the data-collection step, the airline with perceived model i realize m show demands $\{s_{ijt} : j = 1,...,m\}$, the random sample from the actual distribution [the binomial distribution with parameter $x_{it}^*(y_{it})$ and θ]. The realized show-up rate are $\{r_{ijt} : j = 1,...,m\}$ where $r_{ijt} = \dfrac{s_{ijt}}{x_{it}^*(y_{it})}$. In the forecasting step, the show-up rate for the next decision period is forecasted based on the simple exponential smoothing technique: $y_{i,t+1} = \gamma_i \hat{r}_{it} + (1 - \gamma_i)y_{it}$ where $\gamma_i \in (0, 1)$ is the smoothing parameter, and $\hat{r}_{it}$ is the maximum likelihood estimator (MLE) of the parameters of the show-up rate distribution, determined from the realized show-up rate $\{r_{ijt} : j = 1,...,m\}$. For instance, with perceived model 1, $\hat{r}_{1t} = \bar{r}_{1t}$, with perceived

model 2, $\hat{r}_{2t} = \left(\bar{r}_{2t}, \left(\dfrac{\sum_{j=1}^{m}(r_{2jt} - \bar{r}_{2t})^2}{m} \right) \right)$ the MLEs of the parameter (μ, σ^2) of

normal show-up rate distribution, where $\bar{r}_{1t} = \dfrac{\sum_{j=1}^{m} r_{1jt}}{m}$. We restrict our attention to the exponential smoothing method, because of its simplicity and popularity in practice (Makridakis et al. 1998 p. 514), and to the maximum likelihood estimator, because it is one of the most widely used methods of estimation in statistics (DeGroot and Schervish 2002 p. 355). After obtaining the new forecast $y_{i,t+1}$, the process continue with the optimization step in period $(t+1)$ to determine $x_{i,t+1}^*(y_{i,t+1})$, and so on.

4 Numerical Experiment

In this section, two sets of numerical experiments are conducted. In the first set, we study the asymptotic behavior of the perceived model as the number of decision periods in the iterative process becomes very large. In the second set, we compare the per-flight costs if the airline implements the overbooking level from each of the three perceived models. To estimate the expected costs, we perform a Monte Carlo simulation. With the perceived models, the parameters of the show-up rate distribution are updated every $m = 30$ flights, and the smoothing parameters are $\gamma_1 = \gamma_2 = \gamma_3 = 0.5$.

4.1 Asymptotic Behavior Investigate

Data for our numerical experiments are obtained from one of the leading passenger airlines in Thailand. We consider a single-leg weekly flight with capacity $c = 338$ seats. The airline sets the per-unit oversale and spoilage costs to

$a_s = a_o = 4800$, which is the reference fare of the flight. With the actual model, given the overbooking level x the show demand follows the binomial distribution with parameters x and $\theta = 0.945$.

We report only the study of the deterministic model (perceived model 1), because the asymptotic behaviors of the other perceived models look similar. Figure 1 shows four samples paths, when the initial show-up rates are 0.945, 0.945, 0.875 and 0.875, respectively (as indicated in the legend of the figure). From Figure 1a (see appendix), the sequence of the overbooking levels does not converge. For instance, when the initial show-up rate is $\rho_1 = 0.945$, the last five overbooking levels of the first sample path (the solid line) are 357, 357, 357, 357, 358, whereas those of the second sample path (the dotted line) are 357, 358, 358, 357, 358.

Figure 1b (see appendix) suggests that two sample paths of the average costs corresponding to $\rho_1 = 0.945$ converge, and so do the other two corresponding to $\rho_2 = 0.875$. Moreover, all four sample paths of the average costs converge to a single number, which is approximately 16600. With different show-up rates, the sequences of the average costs do converge to the same point. When many replications are carried out, the figure (not shown) suggests that almost all sample paths converge to a single point. We conclude that the long-run average cost converges to a constant with probability one. Nevertheless, it does not converge to the optimal cost based on the actual model, which is 16522.26. Hence, the perceived model is not robust to the misspecification error.

4.2 Performance Evaluation

In the second set of experiments, let $m = 30$, and $c = 338$ (as in the first set). Let the initial forecasts for the three perceived models be $\rho_1 = 0.945$, $(\mu_1, \sigma_1^2) = (0.945, 0.026)$ and $(a_1, b_1) = (62.327, 3.855)$. The rest of the problem parameters, namely the show-up probability θ and the pair of oversale and spoilage costs (a_s, a_o), are varied systematically according to the 3x1 factorial experiment. The first factor-the show-up probability - has 3 levels $\theta \in \{0.8, 0.5, 0.3\}$ and the second factor-the cost parameters-has 1 levels $(a_s, a_o) \in \{(4800, 4800)\}$. The problem parameters for all three experiments are shown in Table 1a (see appendix).

In each simulation replication, we fix the number of decision periods to be 200. (From Figure 1b, the average cost for a given initial show-up rate has already settled down since the 200-th decision period.) The number of simulation replications is chosen such that the length of the 95% confidence interval is within 10% of the estimated cost. Table 1b (see appendix) shows the optimal expected cost based on the actual model, the estimated cost when the airline uses the overbooking level from the perceived model, the corresponding standard error, and the percent difference between the estimated cost and the optimal expected cost.

In each of the three experiments, the estimated cost with the deterministic model (perceived model 1) is highest among the three perceived models. The percent difference with perceived model 1 ranges from 7 to 25. Recall that the deterministic model takes into account only the forecasted show-up rate, whereas the other perceived models take into account not only the forecast but also the cost parameters. Thus, our result that the deterministic model performs

worst is not surprising. Nevertheless, when the show-up probability is high ($\theta = 0.8$), the largest percent difference is no more than 7, which might not be really bad from some RM-industry viewpoints. This together with its simplicity appeals to some RM practitioners.

In each of the three experiments, the estimated cost with the beta model (perceived model 3) is lower than that with the normal model (perceived model 2). With the beta distribution, the percent diference ranges from 4 to 13, whereas it ranges from 6 to 25 with the normal distribution. From our experiences in running these numerical experiments, the computational times of both models are not much different. Hence, the beta show-up rate might be preferred to the normal show-up rate, since it yields a lower expected cost.

From Table 1b, the percent diference increases as the show-up probability decreases. For instance, with $a_s = a_o = 4800$ and with perceived model 3, the percent diference increases from 4 to 9 to 13, when the show-up probability decreases from 0.8 to 0.5 to 0.3. If the airline anticipates a high show-up probability, then making an overbooking decision with the perceived model might be acceptable; however, if it anticipates a low show-up probability, the airline needs to be very cautious using the overbooking level recommended by the perceived model, since the performance of the perceived model gets worse when the show-up probability decreases.

5 Concluding Remarks

Overbooking is one of the core components of passenger airline RM. Major airlines employ commercial RM systems to assist them in making overbooking decisions. According to several journal articles, some overbooking models in practice assume that the show demand is the product of the overbooking level and the show-up rate. However, it follows from probability theory that the random show demand follows a binomial distribution, when each reservation shows up independently and with the same probability. The product form specified in the airline's decision model is incorrect, and a model misspecification occurs. In this article, we explore the consequences of the modeling error.

We use Monte Carlo simulation to estimate the per-flight cost when the perceived model is employed. We consider three show-up rate distributions, namely normal, beta, and degenerate. The show-up rate parameters are periodically updated from the historical records. From the first set of the experiments, we find that the long-run average cost with the deterministic model converges to a single point with probability one, regardless of the initial show-up rates. In the second set of the experiments, the percent diferences from the optimal solution range from 4 to 45. The smallest diference corresponds to the beta model, whereas the largest to the deterministic model. Furthermore, the percent difference increases as the show-up probability decreases.

There are several possible extensions to this article. For example, how often should the airline update the show-up rate parameters? Which forecasting technique would be robust to the model misspecification? If the airline is endowed with the specification test, how long would it take to detect the modeling error? We hope to explore these and other related questions in the future.

Appendix

Proof of proposition 1. It follows from pp. 640-641 in Talluri and van Ryzin (2004) that $E[(S_0(x)-c)^+]$ is convex in x. Recall that a_s, $a_0 \geq 0$. Then, $f_0(x)$ is convex in x; see Bazaraa et al. (2006) p.148. Thus, $x_0^* = \text{argmax } \{x \in \mathbb{N} : f_0(x)-f_0(x-1) \leq 0\}$ $\square$

Proof of proposition 2. With perceived model 1, the objective function in (2) becomes

$$f_1(x \mid \rho) = \begin{cases} -a_s \rho x & \text{if } x \leq \frac{c}{\rho} \\ a_0 \rho x - (a_s + a_o)c & \text{if } x > \frac{c}{\rho} \end{cases}$$

It is strictly decreasing on $(-\infty, \frac{c}{\rho}]$ and strictly increasing on $(\frac{c}{\rho}, \infty)$. Clearly, a minimum occurs at the kink $\frac{c}{\rho}$. For each $i \in \{2, 3\}$, similar arguments as in the proof of Proposition 1 reveal that $f_i(x \mid y_i)$ is convex in x. The optimality condition states that the first derivative is equal to zero. After applying Leibniz' rule and some algebraic simplifications, we obtain (4). $\square$

Figure 1a
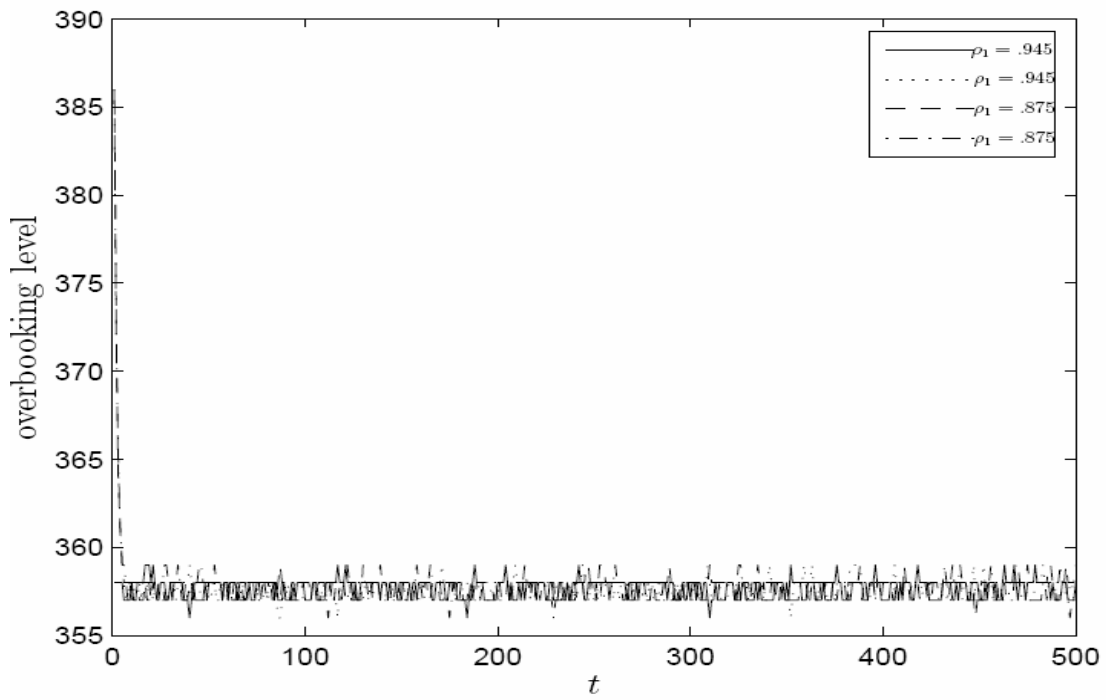
(a) Overbooking levels

Figure 1b
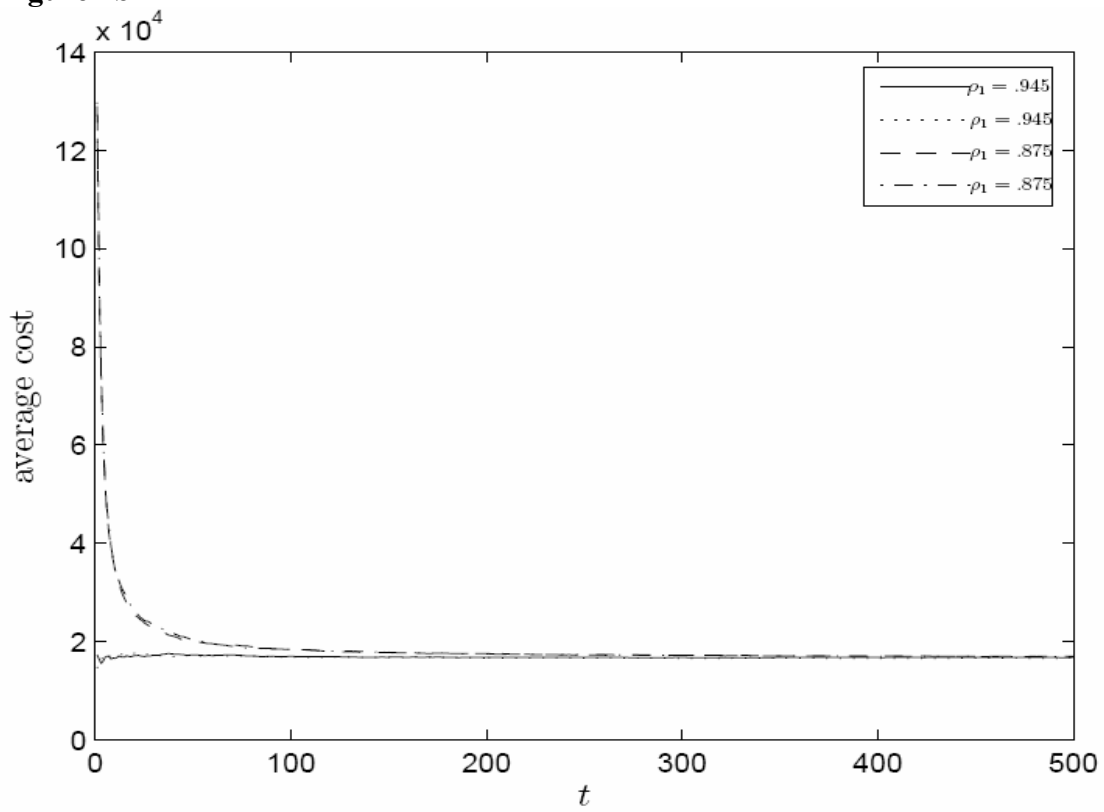
(b) Average costs
Figure 1: Some sample paths with perceived model 1

11

Table 1a

$\theta \backslash (a_s, a_o)$	**(4800,4800)**
0.8	Ex.1
0.5	Ex.2
0.3	Ex.3

(a) Problem parameters for all three experiments

Table 1b

Ex.	Act. Mod. **Opt. Cost.**	Perceived model 1			Perceived model 2			Perceived model 3		
		Mean	%Dif	SE	Mean	%Dif	SE	Mean	%Dif	SE
1	31456.0	33513.2	6.5	9.2	33527.2	6.6	9.5	32848.9	4.4	12.3
2	49769.6	57927.9	16.4	15.9	57937.5	16.4	15.8	54448.9	9.4	52.3
3	58884.4	73317.0	24.5	19.3	73382.2	24.6	19.4	66448.8	12.8	93.2

(b) Costs associated with overbooking levels from perceived models

Table 1: Performance of perceived models

Reference

J. Bailey. Bumped fliers and no plan B. *The New York Times*, 2007. May 30, 2007.

M. S. Bazaraa, H. D. Sherali, and C. M. Shetty. *Nonlinear Programming: Theory and Algorithms*. John Wiley & Sons, Inc., New York, 2006.

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

R. E. Chatwin. Continuous-time airline overbooking with time dependent fares and refunds. *Transportation Science*, 33:182-191, 1992.

W. L. Cooper, T. Homem-de-Mello, and A. J. Kleywegt. Models of the spiral-down effect in revenue management. *Operations Research*, 54(5):968-987, 2006.

J. Coughlan. Airline overbooking in the multi-class case. *Journal of the Operational Research Society*, 50(11):1098-1103, 1999.

M. H. DeGroot and M. J. Schervish. *Probability and Statistics*. Addison-Wesley, Boston, 2002.

G. W. Evans and S. Honkapohja. *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, New Jersey, 2001.

L. P. Hansen, P. Maenhout, A. Rustichini, T. J. Sargent, and M. M. Siniscalchi. Introduction to model uncertainty and robustness. *Journal of Economic Theory*, 128(1): 1-3, 2006.

J. A. Hausman. Specification tests in econometrics. *Econometrica*, 46(6):1251-1271, 1978.

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

R. G. Kasilingam. Air cargo revenue management: Characteristics and complexities. *European Journal of Operational Research*, 96(1):36-44, 1996.

R.G. Kasilingam. An economic model for air cargo overbooking under stochastic capacity. *Computers and Industrial Engineering*, 32(1):221-226, 1997.

R. Klophaus and S. Polt. Airline overbooking with dynamic spoilage costs. *Journal of Revenue and Pricing Management*, 6(1):9-18, 2007.

S. Luo, M. C̜akanyildrim, and R. G. Kasilingam. Two-dimensional cargo overbooking models. *European Journal of Operational Research*, 197(3):862-883, 2009.

S. Makridakis, S. C. Wheelwright, and R. J. Hyndman. *Forecasting:Methods and Applications*. John Wiley & Sons, Inc., New York, 1998.

R. Phillips. *Pricing and Revenue Optimization*. Stanford University Press, Stanford, CA, 2005.

A. Popescu, P. Keskinocak, E. Johnson, M. LaDue, and R. Kasilingam. Estimating air-cargo overbooking based on a discrete show-up-rate distribution. *Interfaces*, 36(3): 248-258, 2006.

B. Presnell and D. D. Boos. The IOS test for model misspecification. *Journal of the American Statistical Association*, 99(465):216-227, 2004.

E. Shlifer and Y. Vardi. An airline overbooking policy. *Transportation Science*, 9:101-114, 1975.

B. C. Smith, J. F. Leimkuhler, and R. M. Darrow. Yield management at american airlines. *Interfaces*, 22(1):8-31, 1992.

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

K. Talluri and G. J. van Ryzin. *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers, Boston, Massachusetts, 2004.

H. R. Thompson. Statistical problems in airline reservation control. *Operations Research Quarterly*, 12:167-185, 1961.

D.-M. Wu. Alternative tests of independence between stochastic regressors and disturbances. *Econometrica*, 41(4):733-750, 1973.