ORiON, Vol. 7, No. 2, pp. 71-94 ISSN 0259-191-X

http://orion.journals.ac.za/

AIRLINE SEAT INVENTORY CONTROL [1] BENEFITING FROM CURRENCY DIFFERENTIALS TO ENHANCE REVENUES

C. E. Love*, W. W. Bird**, R. J. Horsfall† & J. Kingwill‡

* Visiting Professor, Graduate School of Business
University of Cape Town, Rondebosch, 7700
Faculty of Business Administration
Simon Fraser University, Burnaby, Canada

** Civil Engineering Department
University of Stellenbosch, Stellenbosch, 7600

† MBA Candidate, Graduate School of Business
University of Cape Town, Rondebosch, 7700

‡ Civil Engineering Department,
University of Cape Town, Rondebosch, 7700

ABSTRACT

The purpose of this paper is to develop an airline seat inventory control model which will capitalize on currency differentials that exist between city pairs. The approach taken here is to maximize Expected Marginal Seat Revenues as proposed by Belobaba for non-nested fare classes. The basic Expected Marginal Seat Revenue model is extended to explicitly include the effects of overbooking. Data from the South African Airways return flight between CapeTown and London is utilized to demonstrate the model.

[1] This research is supported by NSERC Grant # 1228. The author would like to thank an anonymous referee for many useful comments on an earlier draft of this paper.

1. Introduction

In the post-deregulation airline environment, all carriers have struggled to contain

costs in the face of strong downward pressure on airline fares. Many carriers have also been successful in developing increasingly sophisticated revenue control and yield management techniques. Essentially these techniques have focused on improving airline revenues by monitoring and adjusting the balance between passenger demand and the supply of available seats. For a comprehensive review of the yield management problem refer to Belobaba [2].

While the fare charged is an important aspect of yield management, for the most part, the fare between city pairs is set by competitive forces. Seat inventory control allows an airline to improve yields and total revenues on a flight-by-flight basis within a given fare structure. The major problem is that of determining booking policies for the various classes that will generate the maximum revenue on each flight.

Another important aspect of seat inventory control is to determine an optimal level of overbooking for each class on the flight. Overbooking targets are set to achieve the optimal balance between the benefits of load factor improvements and the cost to the carriers of denied boardings in the event that the number of "no-shows" is less than anticipated. If it is properly controlled, an overbooking policy can benefit both the airline and its customers (Shlifer [11]); carrier profitability is enhanced and at the same time, more of the real customer demand can be satisfied.

In this research we wish to look at the utilization of yield management, and more specifically, the seat inventory control problem when a significant currency differential exists between city pairs. That is, if the rights to the airline seat can be purchased on either end of the city pair, it may be advantageous for a carrier to restrict seat sales at one point in order to gain from (expected) sales at the other point. We see that we have two classes of customers; a high fare class originating out of one city and a second low fare class originating out of the second city.

The literature has addressed the seat inventory control problem essentially from two different frameworks. The first approach has been to allocate the seats into two separate or non-nested groups. The control problem is to determine a rationing scheme to allocate the available seats on the aircraft. The second approach has been to treat the separate classes in a nested fashion. A nested seat sale would give preemptive rights to one of the two cities for any seats available. With a high fare city and a low fare city this would presumably mean that any available seat could be sold to the high fare city. However, in order to ensure that not all the seats are sold to high fare customers, the control problem is to establish a reservation limit for the high fare customers. In this way, yield from the flight can be maximized. A good overview for these two strategies can be found in Belobaba [2].

In this research, we model the problem as a two-class non-nested case. In approaching the problem in this manner, we can incorporate both the inherent co-variability of demand between the two cities and determine an optimal level of overbooking for each flight. While some valuable work has been done in analyzing nested class structures (see Brumelle [5] and [6]), the scenarios being analyzed are too restrictive for use in this situation. Specifically, the basic structure of the nested seat sale is restricted to non-correlated demands in the various classes. Work done on correlated demand in nested structures requires that (beginning with the lowest class) sales in each class must be successively closed off prior to the sales in any higher classes. This is because in nested two-class structures we are determining only one control limit with the higher class having a preemptive right to all the available unsold seats. This is obviously a highly unrealistic assumption for our two city seat sale problem. Using a non-nested analysis, since we are solving for both control limits simultaneously and the preemptive right does not exist, the structure can be disaggregated into two separate decisions. As noted by Belobaba [3], repeated solution of the non-nested model as each available seat is demanded yields the same result as the nested model. Consequently, even assuming that an airline wished to structure their seats in a nested fashion, the non-nested model being employed here will yield good results. Notwithstanding, a good deal of work remains to be done on analyzing nested structures more thoroughly.

The layout of this paper is as follows. In Section 2 we detail the Expected Marginal Seat Revenue (EMSR) model for the non-nested two class case of interest in this situation. In Section 3 we apply the EMSR model to the case of point-of-sale control. In Section 4 we extend the analysis to include the opportunity for overbooking at both points of sale. Finally in Section 5 we present examples of the optimal seat inventory control using data from the CapeTown/London sector of South African Airways.

2. Expected Marginal Seat Revenue Model (EMSR)

The origin of this model was the two class model proposed by Littlewood [10] which categorized passengers as ***high yield*** or ***low yield***. In the model it was recognized that acceptance of a low yield passenger resulting in a subsequent rejection of a high yield passenger meant that the airline lost revenue. This model required that low yield passengers be accepted only if the **certain** revenue exceeded the **expected** revenue of a potential high yield passenger. Belobaba made the distinction between nested and non-nested and developed a model which would incorporate probabilistic demand and multiple fare classes on a single leg flight.

The probability density function for the total number of requests for reservations, r_i received by the airline in fare class i by the close of booking for a particular flight is defined as $p_i(r_i)$. This probability function can be obtained from historical data for the same or similar flights. The cumulative probability that all requests for a particular fare class will be accepted can then be defined as a continuous function of the number of seats S_i allocated to the class:

$$\begin{aligned} P_i(S_i) &= P_r[r_i \leq S_i] = \int_0^{S_i} p_i(r_i)dr_i \text{ and} \\ P_i'(S_i) &= P_r[r_i > S_i] = \int_{S_i}^{\infty} p_i(r_i)dr_i \end{aligned} \quad (1)$$

If we define a booking in class i as b_i, then the expected number of bookings in class i, given a seat allocation S_i, is therefore

$$\bar{b}_i(S_i) = \int_0^{S_i} r_i p_i(r_i)dr_i + S_i P_i'(S_i). \quad (2)$$

It is easy to see that the probability that any given request in class i will be refused is defined as

$$PRR_i = 1 - \frac{\bar{b}_i(S_i)}{\bar{r}_i} \quad (3)$$

where $\bar{r}_i$ is the mean demand in class i.

This probability is to be distinguished from $P_i'(S_i)$ which is the probability that the fare class will be sold out and at least one fare request will be refused.

The aim of the model is to maximize the total expected revenue over all classes. The contribution to the total revenue from each class is given by

$$\overline{R}_i = \overline{R}_i(S_i) = f_i \overline{b}_i(S_i) \tag{4}$$

where f_i is the average fare received from passengers in class i. It follows that the total for a multiclass situation is

$$\overline{R} = \Sigma_i \overline{R}_i. \tag{5}$$

At the optimal point, the expected marginal seat revenue for each class defined is as

$$EMSR_i(S_i) = \frac{\partial \overline{R}}{\partial S_i} = f_i P_i'(S_i) \tag{6}$$

will be equal across all classes. In the case of a two class situation

$$EMSR_1(S_1) = EMSR_2(S_2) \tag{7}$$

for the optimal values of S_1 and S_2.

Equation (7) implies that

$$\frac{P'(S_1)}{P'(S_2)} = \frac{f_2}{f_1}.$$

The interpretation here is clear. The optimal allocation of seats across the two classes is such that the ratio of the probabilities of denied boarding in the two classes equals the reciprocal of the fares of the two classes. That is, the marginal opportunity cost of a denied seat must be the same across the two classes.

The relationships between S_i, $\overline{R}_i$ and C are shown schematically in Figure 1 for the two class case. The base represents the capacity C, and the seat allocations S_1 and S_2 have their origins on the right and left hand sides of the base respectively. As S_1 and hence R_1 increase, so the respective values S_2 and R_2 decrease. The upper part of the figure shows the demand curves with r_1 and r_2 coincident with S_1 and S_2.

3. Extension of the EMSR model to include point-of-sale control

The specific problem under investigation here is the Cape Town/London sector, flown by South African Airlines. Except for a small percentage, passengers travelling on this sector do so on return air tickets. Since First and Business fare classes are not restricted in their length of stay or time of travel, it is reasonable to assume that, given the high frequency of flights, passengers leaving from either Cape Town or London return to their point of departure in some distributed fashion, beginning at some time after their time of departure. That is, for a flight departing from Cape Town for London, there will be claims on seats from these same passengers for the return leg of the flight (London to Cape Town). We assume here that an equilibrium is established such that seat sales from either sale point effectively result in a contemporaneous claim on a return seat. This assumption is substantiated by peak seasonal demand data on flights on this sector. (Horsfall [7]). Consequently, this allows us to approximate each flight, whichever its point of departure, as being made up of customers from two points of sales.

Our problem is to determine how many customers from each sale point to allow on the flight. As we illustrate below in Section 5 using data from the Cape Town/London sector, it is evident from the fare breakdown that there exists an incentive to capture the benefits from differential exchange rates in making the booking decision. Within any travel class, in terms of revenue generation, there are two non-nested fare classes; one for each sale point. It is non-nested since, by virtue of our assumption regarding the contemporaneous nature of the sales, neither sale point has a claim on the seats of the other sale point (Belobaba [2], Brumelle [5]). Hence, it is only necessary to apportion the seats available in the travel class between the two sale points such that the total revenues (for the class) are maximized.

In our numerical examples below, we consider examples from both First class and from Business class travel. We are less confident about the use of this model in the Economy class. This is because the First and Business class can reasonably be mod-

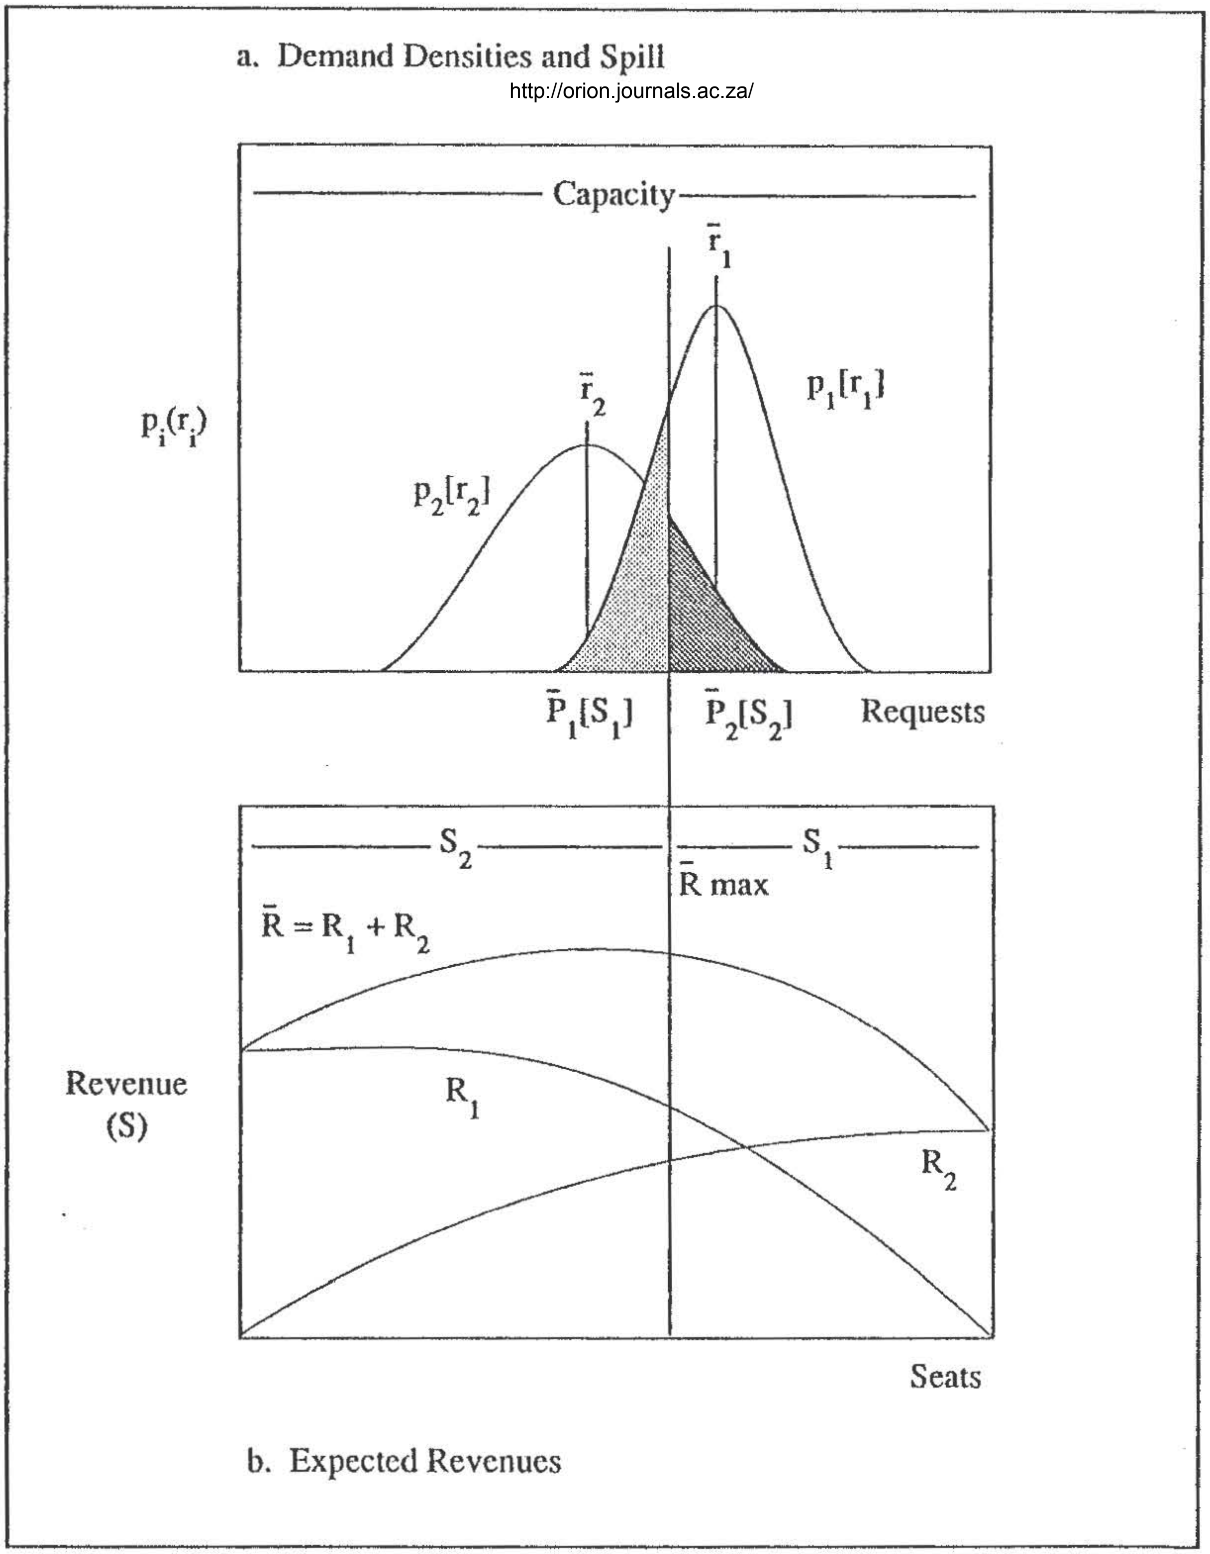

Figure 1: Optimal seat allotment

Source: Belobaba [3], page 106

elled as separate (non-nested) buckets. Each of these two classes has a separate cabin space allocated to it. In the non-nested structure, we can determine the rationing scheme for the seats to be allocated to each city. In addition, given cabin limitations, it is easy to determine the effect of overbooking on the class.

In contrast, modern airline seat management places economy seats within a nested structure (Brumelle [5]). Reservation limits are placed on these higher classes to avoid early, uneconomic consumption of the higher revenue First and Business class seats. The use of the non-nested approach will certainly generate approximate solutions to the Economy class situation, but much more work needs to be done on the two-city nested case. Note that none of these models explicitly treat the upgrading phenomenon prevalent in the airline industry. Here seats reserved for higher classes are "given back" to satisfy unfilled lower class demand. Typically this happens close to flight time when it is discovered that higher class reservations will not all be taken up while some lower lass demand will be left unsatisfied. To the extent that airlines utilize a seat inventory model, however, the upgrading problem with or without nesting can be treated by rerunning the model at various points in time as new demand information is obtained.

Define the expected revenues for travel class j as:

$$\overline{R}^j = \overline{R}_1^j + \overline{R}_2^j \tag{8}$$

where the subscripts 1 and 2 refer to the two separate fare classes i.e. either London or Cape Town generated revenue. Since we will only be considering one travel class at a time, the superscript denoting the travel class will be dropped for the remainder of this discussion. For fares of f_1 and f_2 in fare classes 1 and 2 respectively, the expected revenue for the travel class can be expressed as

$$\overline{R} = f_1 E(r_1/S_1) + f_2 E(r_2/S_2). \tag{9}$$

We will concentrate on the first term on the right hand side of the above equation, namely $\overline{R}_1$. This term is represents the expected number of reservations, r_1 in fare class 1 given the number of seats, S_1 allocated to the fare class multiplied by the fare. In expanded form, this revenue is

$$\overline{R}_1 = f_1\{\int_0^{S_1} r_1 p(r_1)dr_1 + S_1\int_{S_1}^{\infty} p(r_1)dr_1\} = f_1\overline{b}(S_1). \tag{10}$$

Assuming normally distributed demands and letting

$$z_1 = \frac{r_1 - \mu_1}{\sigma_1},$$

we can write (10) as

$$\overline{R}_1 = f_1[\frac{1}{\sqrt{2\pi}}\int_{-\frac{\mu_1}{\sigma_1}}^{\xi_1}(\mu_1 + \sigma_1 z_1)e^{-z_1^2/2}dz_1 + \frac{1}{\sqrt{2\pi}}S_1\int_{\xi_1}^{\infty} e^{-z_1^2/2}dz_1] \tag{11}$$

where the transformation,

$$\xi_1 = \frac{S_1 - \mu_1}{\sigma_1}$$

is applied to the limits of the integral.

Integrating (11), it is easy to establish that

$$\overline{R}_1 = f_1\{\mu_1[\Phi(\xi_1) - \Phi(-\frac{\mu_1}{\sigma_1})] - \sigma_1[\phi(\xi_1) - \phi(-\frac{\mu_1}{\sigma_1})] + S_1[1 - \Phi(\xi_1)]\} \tag{12}$$

where Φ is the standard normal distribution function and ϕ is the standard normal density function.

Φ can be evaluated by means of a numerical method such as that proposed by Abramowitz and Stegun [1]. The remainder of the terms can be computed explicitly. The same argument leads to the revenue $\overline{R}_2$ derived from fare class 2 and hence the total revenue for the travel class can be computed. Substituting (11) and the equivalent result for $\overline{R}_2$ into equation (9) and using the fact that with no overbooking allowed $S_1 = C - S_2$, optimal conditions can be found by solving:

$$\frac{d\overline{R}}{dS_1} = 0. \tag{13}$$

The solution of (13) yields the result:

$$f_1[1 - \Phi(\frac{S_1 - \mu_1}{\sigma_1})] = f_2[1 - \Phi(\frac{S_2 - \mu_2}{\sigma_2})] \qquad (14)$$

which is a simple restatement of equation (7). Consequently it is easy to see that a simple line search on S_1 (or S_2) will determine the optimum values of S_1 (and S_2) in the overbooking situation.

4. The Overbooking Problem

The major benefit achieved by increasing the seat allocation in each fare class is that the overall load factor increases (i.e. the average number of passengers per flight) and hence the expected revenue. The drawback is that with the increased seat allocation, there is the possibility of having to reject passengers resulting in a loss of the fare for the flight plus additional expenses associated with denied boarding. In the data we supply below, it is evident that these additional expenses are non-trivial.

In considering the overbooking problem here, we do not include the possibility of post-purchase cancellations as in Shlifer [11] or Brumelle [6]. That is, in our analysis we assume that a seat once booked is also confirmed, whereas in Shlifer [11] and Brumelle [6] the assumption is that one is dealing with a two stage process; customers first book and then elect to confirm or cancel (as a Bernoulli process). Our cancellation process is implicit in the variance of the demand estimates and our sole concern is with an attempt to increase the load factor balanced against the possibility of additional overbooking charges when we have overfilled the travel class. In our approach, we could accommodate post-purchase cancellations in a more explicit manner either by scaling down the real demand by some percentage or by reducing the overbooking charges in an equivalent way. The net effect would be to overbook more in order to adjust for the fact that not all customers will actually show up.

The impact of the expected lost revenue from overbooking is not trivial, as the costs associated with denied boarding depend on the relative overbooking costs at the two departure sites (in our case, London overbooking charges are significantly higher than Cape Town). We will consider two approaches to this problem. In the first and simpler case, it is assumed that the overbooking costs in both Cape Town and London are the same. This automatically leads to the model favouring London seats since they generate the most revenue. This is clearly a shortcoming, since the overbooking costs associated with denied boarding in London are much higher. The second model attempts to address this shortcoming by apportioning the overbooked seats to the two sale points on some rational basis.

Model 1

In order to describe the model, it is necessary to define some additional terms. Let B_1 and B_2 be the booking limits for fare classes 1 and 2 respectively. The sum of the limits exceeds the capacity of the travel class by the extent of overbooking allowed. Mathematically these relationships are expressed as $B = B_1 + B_2$ and $B = C + S^{over}$, where S^{over} is the number of overbooked seats. Hence expected revenues $\overline{R}_1$ for city 1 can be computed directly using B_1 in place of S_1 in equation (12); and likewise a comparable calculation yields $\overline{R}_2$. Since B_1 and B_2 will in general be larger than S_1 and S_2 respectively, total revenues will increase.

To determine the total expected number of rejected seats the two demand distributions need to be summed. The total demand distribution can thus be represented by $N(\mu_3, \sigma_3)$, where the mean and standard deviations are computed as follows:

$$\mu_3 = \mu_1 + \mu_2 \tag{15}$$

and

$$\sigma_3 = \sqrt{\sigma_1^2 + \sigma_2^2 + 2\rho_{12}\sigma_1\sigma_2}, \tag{16}$$

where ρ_{12} is the correlation coefficient between the demand of fare class 1 and fare class 2.

We express the normal probability density function for the sum of the demands as $p(r_3)$. If the revenue lost from a single rejected seat is f_3, the total revenue expected to be lost is therefore

$$\overline{R}_3 = f_3\{\int_C^B (r_3 - C)p(r_3)dr_3 + (B - C)\int_B^\infty p(r_3)dr_3\}. \tag{17}$$

This lost revenue is clearly zero for $B \leq C$. The net expected revenue can thus be rewritten as

$$\overline{R} = \overline{R}_1 + \overline{R}_2 - \overline{R}_3 \tag{18}$$

which is a function of B_1, B_2 and C. Using equations (12) and (17) and defining $\xi_1 = (B_1 - \mu_1)/\sigma_1$, $\xi_2 = (B_2 - \mu_2)/\sigma_2$ and $\xi_3 = (B - \mu_3)/\sigma_3$ we can rewrite (18) as:

$$\begin{aligned} \overline{R} &= f_1\{\mu_1[\Phi(\xi_1) - \Phi(-\frac{\mu_1}{\sigma_1})] - \sigma_1[\phi(\xi_1) - \phi(-\frac{\mu_1}{\sigma_1})] + B_1[1 - \Phi(\xi_1)]\} \qquad (19) \\ &+ f_2\{\mu_2[\Phi(\xi_2) - \Phi(-\frac{\mu_2}{\sigma_2})] - \sigma_2[\phi(\xi_2) - \phi(-\frac{\mu_2}{\sigma_2})] + B_2[1 - \Phi(\xi_2)]\} \\ &- f_3\{\mu_3[\Phi(\xi_3) - \Phi(\frac{C - \mu_3}{\sigma_3})] - \sigma_3[\phi(\xi_3) - \phi(\frac{C - \mu_3}{\sigma_3})] \\ &\quad +B[1 - \Phi(\xi_3] - C[1 - \Phi(\frac{C - \mu_3}{\sigma_3})]\} \end{aligned}$$

From this it is straightforward to establish, for any given level B of overbooking

$$\begin{aligned} \frac{d\overline{R}}{dB_1} &= \frac{\partial \overline{R}}{\partial B_1} + \frac{\partial \overline{R}}{\partial B_2}\frac{dB_2}{dB_1} \qquad (20) \\ &= f_1[1 - \Phi(\xi_1)] - f_2[1 - \Phi(\xi_2)] = 0. \end{aligned}$$

Note that the optimality condition given by equation (20) is the same as that presented in equation (14). This follows from the fact that the overbooking cost for a given B is independent of the seat allocation.

It is easy to conduct a gradient search using the optimality conditions in (20) to determine the optimum booking levels B_1 and B_2 for each level of overbooking B. To obtain the optimum level of overbooking, a simple line search will suffice:

- *Step a:* Choose a total level B.
- *Step b:* Conduct a gradient search to find optimal partition of B into B_1 and B_2. Use equation (19) to evaluate this allocation.
- *Step c:* Choose another booking level B. If our line search exhausts all values of B stop otherwise return to ***step b***.

Model 2

In reality, a denied booking in one site results in a greater loss of revenue than in the other site (in our case London being the more expensive).

Once again we have the booking limits B_1 and B_2 yielding a total booking level $B(= B_1 + B_2)$. To evaluate the overbooking charge we must separate $\overline{R}_3$ into two costs, $\overline{R}_{3_1}$ and $\overline{R}_{3_2}$. $\overline{R}_{3_1}$ represents the cost of overbooking of customers in city 1 who are then denied boarding. $\overline{R}_{3_2}$ is the equivalent overbooking cost for customers booked in city 2. We denote the lost revenue for each seat in the fare classes by f_{3_1} and f_{3_2}. Multiplying these revenues by the expected number of denied boardings for the respective fare classes yields the total expected lost revenue, $\overline{R}_3$.

Consistent with equation (17), denote the expected number of denied bookings to be:

$$E(d_b) = \int_C^B (r_3 - C)p(r_3)dr_3 + (B - C)\int_B^\infty p(r_3)dr_3. \tag{21}$$

We define β_1 as the probability that a random customer on either leg of the flight was booked in city 1. From equation (2) this is,

$$\beta_1 = \frac{\overline{b}_1(B_1)}{\overline{b}_1(B_1) + \overline{b}_2(B_2)}. \tag{22}$$

Since, if we must remove customers from the flight and customers on the return flight will have preemptive right to a seat, the proportion of denied boardings (overbookings) that will take place in city 1 is β_1 and the proportion that will take place in city 2 is β_2. To keep flights from being overloaded, customers must be denied flying out of the two cities in proportions β_1 and β_2 respectively. Hence the total expected cost of overbooking can be computed as

$$\overline{R}_3 = f_{3_1}\beta_1 E(d_b) + f_{3_2}\beta_2 E(d_b). \tag{23}$$

Substituting (23) into (18) it follows that

$$
\begin{aligned}
\frac{d\overline{R}}{dB_1} &= f_1[1-\Phi(\xi_1)] - f_2[1-\Phi(\xi_2)] \qquad (24)\\
&\quad -\{\frac{(\bar{b}_1[1-\Phi(\xi_2)] + \bar{b}_2[1-\Phi(\xi_1)])(f_{3_1} - f_{3_2})}{[\bar{b}_1+\bar{b}_2]^2}\}E(d_b)\\
&= 0
\end{aligned}
$$

which allows us to determine the optimum booking levels B_1 and B_2 in the presence of differential overbooking charges in the two cities. Note that we are making these booking limit decisions, B_1 and B_2, in advance of flight time. In this way, we are balancing off the expected revenues of bookings in the two cities against the expected costs of having too many customers showing up for a flight. The assumption is that customers flying their return leg have a preemptive right to a seat. Only outbound customers may be denied a seat. This does not preclude, however, a returning customer voluntarily accepting the offer to be "bumped".

5. The CapeTown/London Return Sector for South African Airways

Parameter Estimation

Our application of these models is applied to the Cape Town/London sector flown by South African Airways [2]. The method used for estimating the means and standard deviations follows the approach adopted by Boeing Commercial Aircraft Company (Sparham [12]). Normal probability paper is used to plot historical flight data. Capacity used (on the vertical axis) is plotted against frequency of flights using no more than this capacity (on the horizontal axis). If the data follows a normal distribution, then the plot would follow a straight line except for the upper end where the curve would bend down because of flights flown on which cabin capacity was reached. The data used in the examples used in this paper exhibited this behaviour.

To estimate the unconstrained demand, a simple estimation technique is to utilize the 50th percentile (the median) as an estimate of the mean and to utilize the slope of the normal probability plot as an estimate of the standard deviation. This can be seen because between the 16th and the 50th percentiles would constitute 1 standard deviation. Hence $z = 1 = (\mu_{50th} - x_{16th})/\sigma$. From this it is clear that $\sigma = \mu_{50th} - x_{16th}$.

Provided we use the best-fit regression to establish the slope of the normal proba-

[2] The authors wish to express their appreciation to the South African Airways and in particular to Louis du Plessis for many fruitful talks and for supplying the sector data.

bility plot, we can determine our distribution parameters simply and quickly. This does, however, remain censored data (depending on the number of flights flown fully loaded). More efficient (lower variance) parameter estimates can be obtained via Maximum Likelihood Estimation (the details of which are available in Lawless [9]) for normally distributed censored data. Since it is not the intention of this paper to focus on parameter estimation, the simplified probability plot approach was utilized here. For an actual implementation of these algorithms, efficiency of estimation could become important, in which case the Maximum Likelihood approach should be used.

Below we present the results on several different scenarios using the above models. As mentioned, we have restricted our cases to First and Business Class travel where these models seem to most appropriately fit the structure of the decision process. Further development is necessary to extend these models to economy class where nesting is now the dominant structure. In all the examples presented we have assumed that $\rho_{12} = 0$ and have briefly discussed the influence of covariance separately.

Example 1: *First Class travel with various levels of overbooking, Model 1.*

London — Cape Town	Cape Town — London
f_1 = R 17035 $\mu_1 = 22$ $\sigma_1 = 11$	f_2 = R 10262 $\mu_2 = 58$ $\sigma_2 = 17$, $\rho_{12} = 0$
Overbooking cost, f_3 = R 18885	

Table 1: Data for First Class with Common Overbooking Charge

The data and results for this example are tabulated in Tables 1 & 2 respectively. The results are also graphed in Figure 2. The capacity of the cabin is 112. (Note that this is double the physical capacity of the First Class cabin since we are booking the return flight simultaneously.)

While we have limited the sales to a maximum of 133 seats (i.e 19% overbooked) we have not reached a revenue maximizing position. In this case $\overline{R}$ continues to increase and as can be seen from Figure 2, has still not reached an optimum at a 50% level of overbooking. Since this is First class, it is unlikely that even a 10% level would be acceptable, signalling that the implied overbooking cost used by the airline is higher than the real cost utilized here.

B	$\overline{R}$	B_1	$\overline{R}_1$	PRR_1	B_2	$\overline{R}_2$	PRR_2	$\overline{R}_3$
112	949596.6	37	368920.9	.016	75	580675.7	.024	.0
113	950128.4	37	368920.9	.016	76	582232.2	.022	1024.7
114	950621.4	37	368920.9	.016	77	583651.3	.019	1950.7
115	951140.2	38	370274.7	.012	77	583651.3	.019	2785.8
.	.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.	.
123	955142.5	41	373155.5	.004	82	588977.5	.010	6990.4
.	.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.	.
132	958572.1	44	374770.0	.000	88	592490.0	.005	8687.9
133	958855.8	44	374770.0	.000	89	592863.7	.004	8777.8

Table 2: Results for First Class with Common Overbooking Charge

Our second level decision is how to allocate the seats between Cape Town sales and London sales. The optimum is obtained by solving equation (20). Since we do not have an revenue maximising number of seats, we have determined the optimum allocation of seats for the base case of 112 seats (i.e. no ovverbooking). We can see that for this case, 75 seats should be sold out of Cape Town and 37 seats should be sold out of London. In this case the probabilities of arriving customers being refused are respectively 1.6% and 2.4% in London and Cape Town. Without running a controlled case, it is obvious that any arbitrary allocation of seats between the two sale points must be inferior to those generated by the model (by virtue of the optimization).

Example 2: *First Class travel with various levels of overbooking, Model 2.*

London — Cape Town	Cape Town — London
f_1 = R 17035	f_2 = R 10262
$\mu_1 = 22$	$\mu_2 = 58$
$\sigma_1 = 11$	$\sigma_2 = 17$, $\rho_{12} = 0$
Overbooking cost, f_{3_1} = R 18885	Overbooking cost, f_{3_2} = R 11662

Table 3: Data for First Class with Separate Overbooking Charge

Here we have a re-run of example 1, but in this instance model 2 has been used in

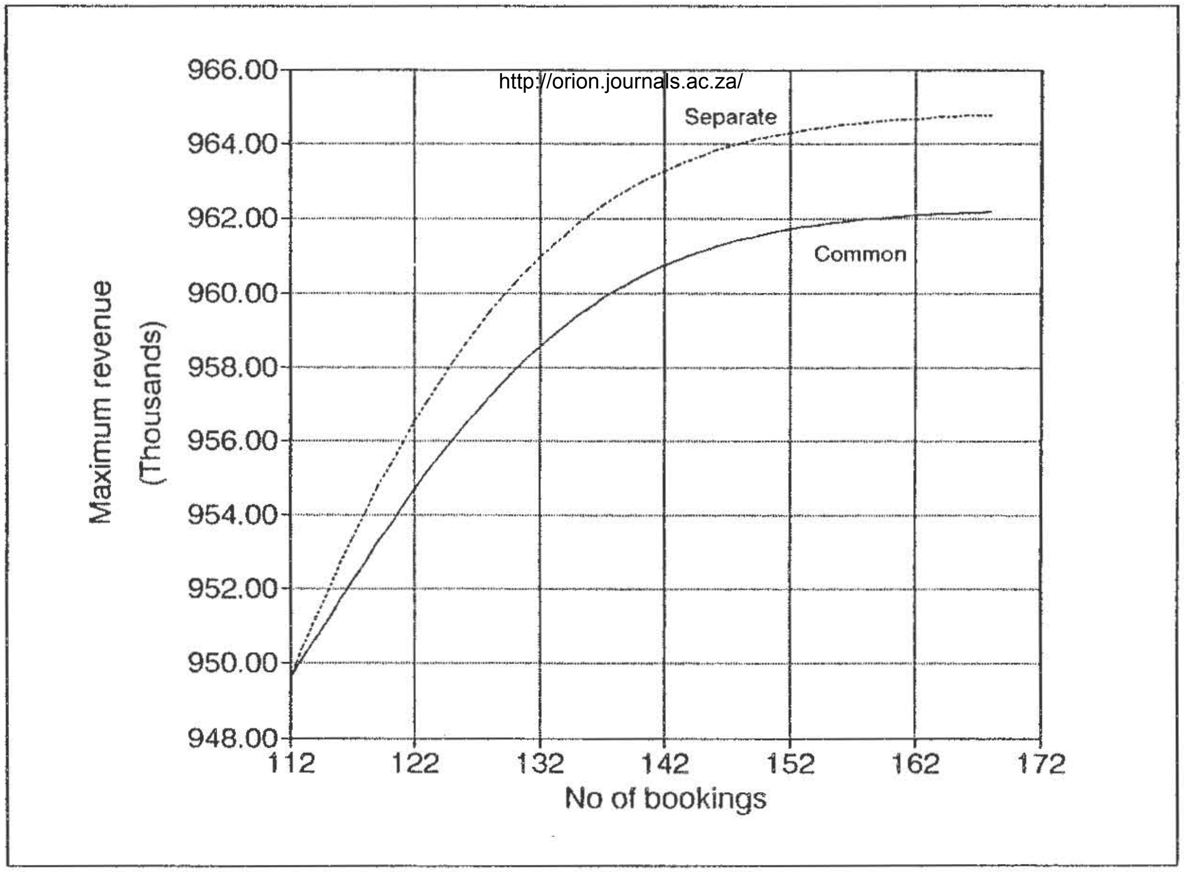

Figure 2: Results for First Class

order to measure the effects of the differentiated overbooking charges. The data used for this example appear in Table 3. The results are tabulated in Table 4 and graphed in Figure 2. As in the previous example no optimum is reached. The expected revenues for different levels of overbooking are slightly higher as the aggregate cost of overbooking in this case is less than in the previous example. If we compare the seat allocations at the 10% level of overbooking (i.e. 123 seats), we see that both models allocate 41 seats to London and 82 to Cape Town. Even though the cost of an overbooked Cape Town seat is less, the likelihood of it occuring is much higher.

Example 3: *J (Business) Class travel, overbooking allowed, Model 1*

The capacity of this class is 176 in both legs. The data used in this example are shown in Table 5. The expected revenues for various overbooking limits are tabulated in Table 6 and plotted in Figure 3 and it is evident that by using this model, the maximum will only be achieved at a very high overbooking level (probably unac-

B	$\overline{R}$	B_1	$\overline{R}_1$	PRR_1	B_2	$\overline{R}_2$	PRR_2	$\overline{R}_3 + \overline{R}_4$
112	949596.6	37	368920.9	.016	75	580675.7	.024	.0
113	950412.1	37	368920.9	.016	76	582232.2	.022	741.0
114	951161.8	37	368920.9	.016	77	583651.3	.019	1410.3
115	951911.1	38	370274.7	.012	77	583651.3	.019	2014.9
.	.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.	.
123	957077.6	41	373155.5	.004	82	588977.5	.010	5055.3
.	.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.	.
132	960978.2	44	374770.0	.000	88	592490.0	.005	6281.8
133	961287.2	44	374770.0	.000	89	592863.7	.004	6346.4

Table 4: Results for First Class with Separate Overbooking Charge

London — Cape Town	Cape Town — London
f_1 = R 9620 $\mu_1 = 49$ $\sigma_1 = 19$	f_2 = R 7280 $\mu_2 = 75$ $\sigma_2 = 33$, $\rho_{12} = 0$
Overbooking cost, f_3 = R 11470	

Table 5: Data for Business Class with Common Overbooking Charge

ceptably high to management). Hence in order to compare this case with example 4 below, we choose a total seat capacity of 200 seats. This represents an overbooking level of 14%. Of these seats, 121 are allocated to Cape Town and 79 are allocated to London with probabilities of arriving customers being refused respectively 1.5% and 0.9%. Note from Figure 3 that the total revenue begins to flatten off at a level of overbooking of about 25%.

Example 4: *J (Business) Class travel, overbooking allowed, Model 2*

Once again this example uses the data (see Table 7) of example 3, but in this instance employing Model 2. As in example 3, a maximum will only be reached at a very high level of overbooking. The results for increasing overbooking limits are tabulated in Table 8 and graphed in Figure 3. Comparing the seat allocations of the two models

B	$\overline{R}$	B_1	$\overline{R}_1$	PRR_1	B_2	$\overline{R}_2$	PRR_2	$\overline{R}_3$
176	983771.6	70	459253.7	.026	106	524517.9	.039	.0
177	984048.5	71	460494.1	.023	106	524517.9	.039	963.4
178	984367.7	71	460494.1	.023	107	525754.8	.037	1881.3
.	.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.	.
200	991709.0	79	467202.5	.009	121	538025.1	.015	13518.6
.	.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.	.
262	1000849.0	101	471494.4	.000	161	546607.6	-.001	17253.0
263	1000879.3	101	471494.4	.000	162	546639.5	-.001	17254.6
264	1000907.0	101	471494.4	.000	163	546668.7	-.001	17256.0

Table 6: Results for Business Class with Common Overbooking Charge

London — Cape Town	Cape Town — London
f_1 = R 9620 $\mu_1 = 49$ $\sigma_1 = 19$	f_2 = R 7280 $\mu_2 = 75$ $\sigma_2 = 33$, $\rho_{12} = 0$
Overbooking cost, f_{3_1} = R 11470	Overbooking cost, f_{3_2} = R7480

Table 7: Data for Business Class with Separate Overbooking Charge

at 200 seats booked, it is evident that model 2 favours Cape Town seats because of the lower overbooking charge. On the whole however, splitting the overbooking cost has little effect in this case and the differences between model 1 and model 2 are not that significant. Once again the slightly lower loss of revenue from overbooking predicted by model 2 is due to the fact that the average overbooking cost per seat is lower.

The influence of covariance

To measure the impact of covariance on the booking limits we have repeated Examples 3 and 4, but with a range of correlation coefficients from -1 to +1. The results of these runs are graphed in Figures 4 and 5. It is clear that for positive correlation a distinct maximum is obtained with the optimum level of overbooking falling off as

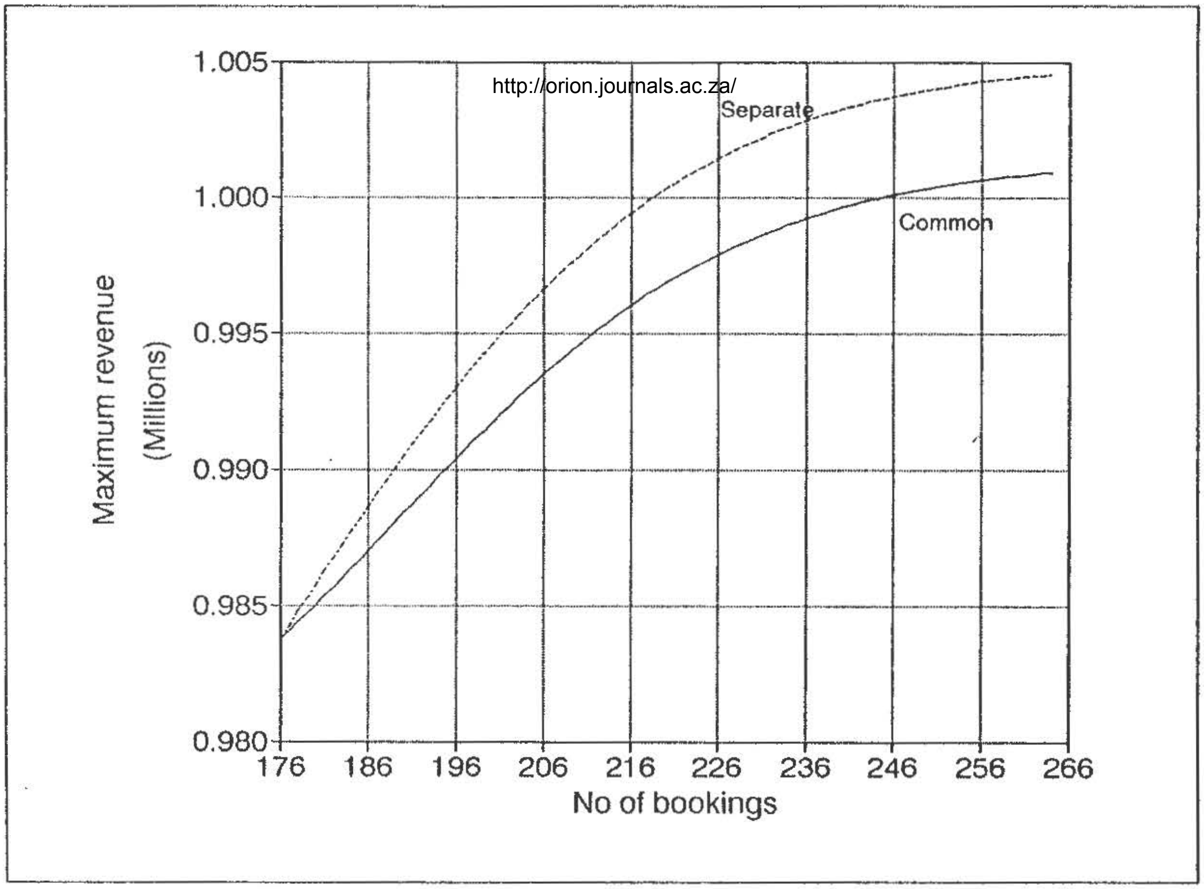

Figure 3: Results for Business Class

ρ_{12} increases.

One should note that the results obtained are for a particular point in time. In reality, the seat allocation problem must be solved dynamically as the time to flight departure (in either direction) occurs. At any given time, the mean and variance of demand for service at both departure points must be estimated. What one observes is that the variability in demand reduces as one approaches the departure date. This is no doubt in part a reflection of the cancellation process in operation. The closer one gets to flight time, the less is the uncertainty that exists in the demand for service. Historically, flight specific data reflects this information in the form of reduced demand variance. In running the model, the smaller the assumed total variance in demand, the lower will be the optimal the level of overbooking (for any given mean demand for service). In this way, the model proposed here can be regarded as a mathematical interpretation of what currently occurs in practice. As flight time is approached an airline reduces the overbooking level. They do this obviously because

B	$\overline{R}$	B_1	$\overline{R}_1$	PRR_1	B_2	$\overline{R}_2$	PRR_2	$\overline{R}_3 + \overline{R}_4$
176	983771.6	70	459253.7	.026	106	524517.9	.039	.0
177	984249.9	71	460494.1	.023	106	524517.9	.039	762.1
178	984761.2	71	460494.1	.023	107	525754.8	.037	1487.7
.	.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.	.
200	994547.7	78	466622.6	.010	122	538603.3	.014	10678.2
.	.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.	.
262	1004480.3	101	471494.4	.000	161	546607.6	-.001	13621.7
263	1004511.0	101	471494.4	.000	162	546639.5	-.001	13622.9
264	1004539.2	101	471494.4	.000	163	546668.7	-.001	13623.9

Table 8: Results for Business Class with Separate Overbooking Charge

the likelihood of cancellation is becoming less. Hence, if they do overbook there is a high probability that they will incur the overbooking charge.

6. Conclusions

We have seen in this research that it is possible for an airline to manage its inventory of seats in such a way as to capitalize on currency differentials that exist between city pairs. While the overall gain is small on any given flight, it is the fact that an airline repeatedly flies such sectors that makes the optimization of value. The fact that only small improvements can be obtained on any one return flight is consistent with previous research in seat inventory management (Brumelle [6])

A good deal of work still needs to be done in this area. To begin with, further calibration of the model needs to be carried out in order to measure the impact on other sectors. Secondly the impact of one-way flights needs to be investigated. This would principally arise in this context, because passengers would fly only one leg of the sector, but elect to return on another airline. The fact that this has a small effect on the sector is what motivated the adopted approach. Finally, the issue of nesting needs to be addressed within the model. This is of course a complex modelling problem, but could potentially be of benefit particularly in the economy class.

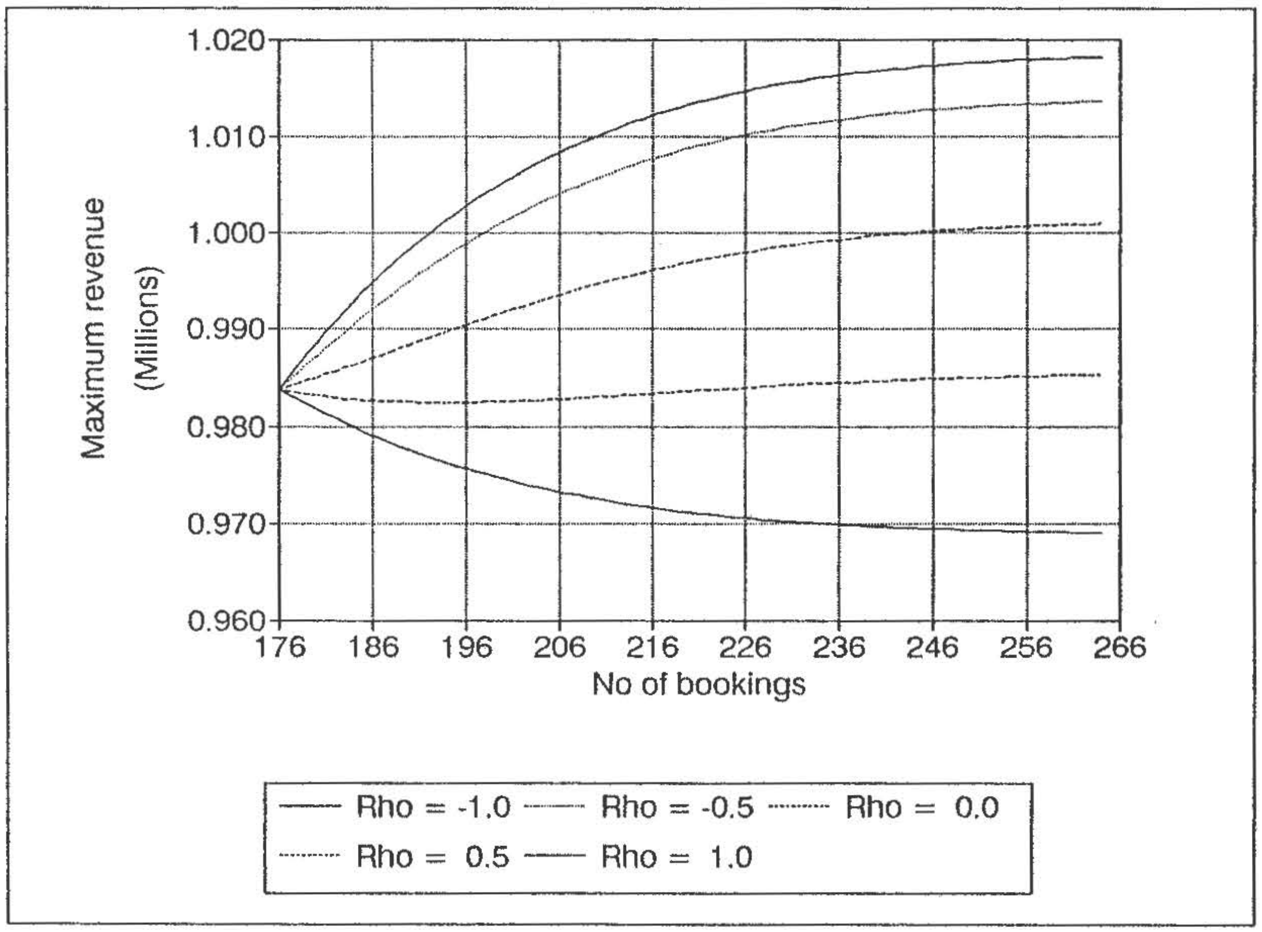

Figure 4: Business Class with Common Overbooking Charge and Varying ρ_{12}

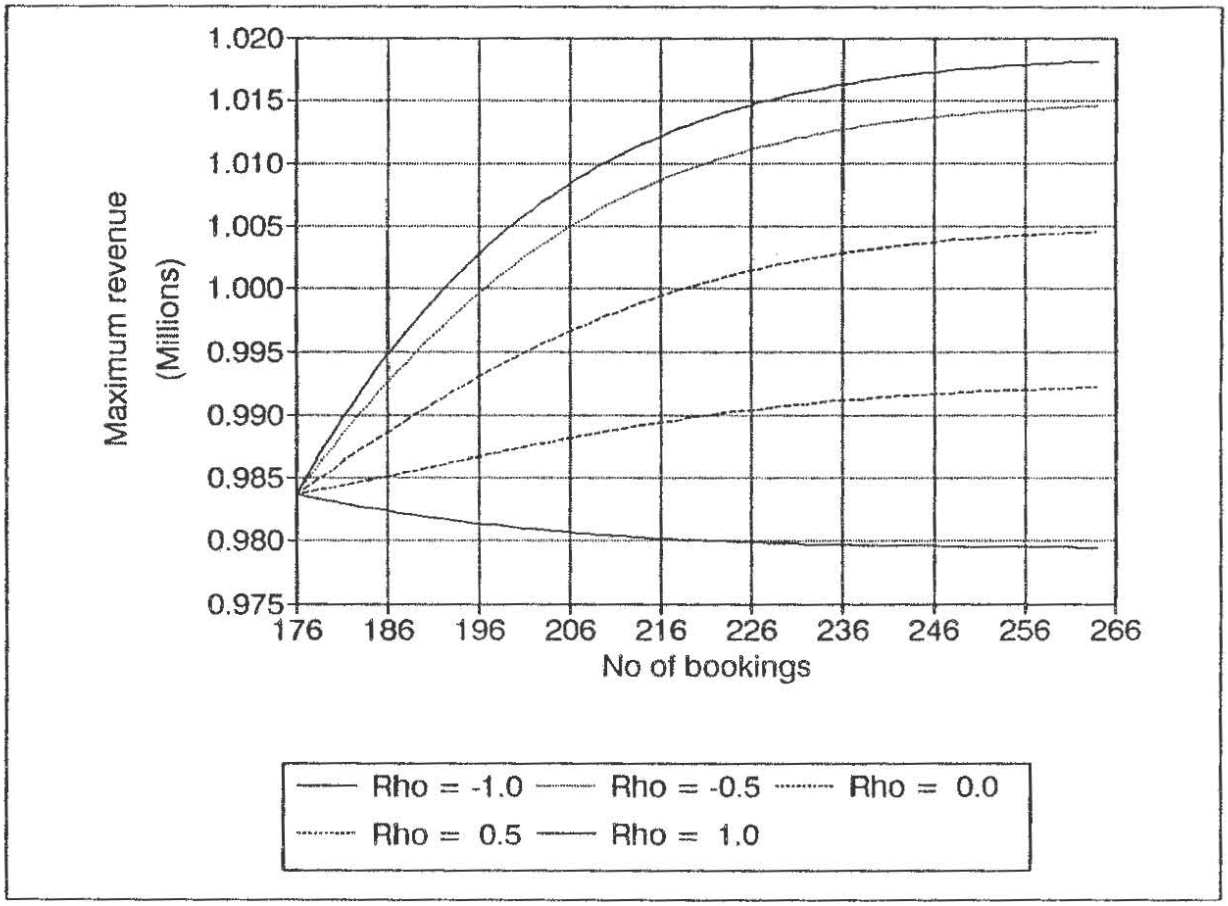

Figure 5: Business Class with Separate Overbooking Charge and Varying ρ_{12}

References

[1] Abramowitz, M. and Stegun, I., *Handbook of Mathematical Functions, NBS Applied Mathematical Series 55*, 9th Printing, Washington DC: National Bureau of Standards. USA., 1970.

[2] Belobaba, Peter P.,"Survey Paper: Airline Yield Management, An Overview of Seat Inventory Control", *Transportation Science*, Vol 21, No 2, pp 63–73, 1987.

[3] Belobaba, Peter P., "Air Travel Demand and Airline Seat Inventory Management," PhD Dissertation, MIT, Cambridge, Mass., 1987, pp. 82—94, 1.

[4] Brenner, M.A., "The Significance of Airline Passenger Load Factors", in G.W. James(ed.), Airline Economics, Lexington Books, Lexington Massachusetts, 1982.

[5] Brumelle, S.L., & McGill, J.I., "Airline Seat Allocation with Multiple Nested Fare Classes", University of British Columbia Working Paper 89-MSC-001. (To appear *Transportation Sciences*), 1989

[6] Brumelle, S.L.,McGill, J.I., Oum, T.H., Sawaki, K., Tretheway, M.W., "Allocation of Airline Seats between Stochastically Dependent Demands",University of British Columbia Working Paper, (To appear, *Transportation Sciences*), 1989.

[7] Horsfall, R.J., "Airline Seat Inventory Control: Benefiting from Currency Differentials to Enhance Revenues", Unpublished MBA Thesis, Graduate School of Business, University of Cape Town, Cape Town, SA, 1990.

[8] Kraft, D.J.H., Oum, T.H. and Tretheway, M.W., "Airline Seat Management", *Logistics and Transportation Review*, Vol 22, No. 2, pp 115–130.

[9] Lawless, J. F., *Statistical Models and Methods for Lifetime Data*, John Wiley & Sons, New York, 1982, pp. 220—237.

[10] Littlewood, K., "Forecasting and Control of Passengers", in Proceedings 12th AGIFORS Symposium, pp 95-117, American Airlines, New York, 1972.

[11] Shlifer, R., and Vardi, Y., "An airline overbooking policy", *Transportation Science*, 9:101–114, 1975.

[12] Sparham, A. H., "Choosing the Optimal Aircraft Configuration: A Yield Management Approach", Unpublished MBA Thesis, Faculty of Business Administration, Simon Fraser University, Canada, 1989.