JOURNAL OF THE INTERNATIONAL
ACADEMY OF HOSPITALITY
RESEARCH

J I A H R

November 26, 1990 ISSN 1052-6099 Issue 1

Editor: Mahmood A. Khan, Virginia Polytechnic Institute & State University
Managing Editor: Eliza Tse, Virginia Polytechnic Institute & State University

- - - - - - - - - - - - -

JIAHR publishes refereed papers on all aspects of hospitality and tourism research. Each issue is largely composed of a single paper and is sent via Bitnet and the Internet to members of the International Academy of Hospitality Research and to subscribing individuals and libraries.

INSIDE

- - - - - - - - - - - - -

Hotel Yield Management Using Optimal Decision Rules

Ralph D. Badinelli
Michael D. Olsen

Virginia Polytechnic Institute and State University
Blacksburg, VA

November, 1990

ABSTRACT

This paper demonstrates a particular model for making the pricing decisions associated with hotel booking. Implementing such pricing decisions that are designed to optimize the profitability of the hotel forms part of a policy commonly referred to as yield management. The model utilizes forecasts of demand in individual market segments to capitalize on the willingness of people in one segment to pay more than people in another segment. The procedure for doing this is necessarily time-based since the market segments are differentiated also by the timing of bookings relative to a rental date. The procedure for making the pricing decisions is described and an example is given. Unlike the commonly invoked marginal revenue models, this model is optimal and requires fewer assumptions about the demand process. It is shown that the procedure has rather modest information requirements and is based on data that is typically available through market research. We also show that the procedure demands minimal amounts of CPU time making it applicable even in small hotels.

INTRODUCTION

The term "yield management" is used to label many approaches to maximizing the profitability of a hotel through manipulation of its pricing and booking policies. The goal of a yield-management system is to consistently maintain the highest possible revenue from a given amount of room capacity. To achieve this goal, the yield-management process includes determining policies for overbooking and allocating hotel capacity to customers of different revenue-generating potential through discriminatory pricing. Ideally, both of these policies should be concurrently incorporated in a hotel’s reservation system. However, it is beyond the scope of this paper to prescribe an optimal policy for simultaneously planning both functions in a coordinated manner. The objective of this paper is to present a model developed for the pricing policy.

The overbooking policy deals with the likelihood of cancellations and no-shows and the consequent lost revenue. Balancing the expected lost revenue due to unoccupied rooms against the loss of goodwill caused by not honoring overbooked reservations is the essential consideration in determining an overbooking policy.

The pricing problem can be identified in two forms. First, there is the revealed-price (RP) case in which a customer

calling in a reservation is assumed to be able to identify his/her customer class, thereby receiving a certain rate. An example of such a situation would be the case of discount rates allowed for attendees to a convention or vacationers who have been given a special group rate. The booking policy that must be determined is usually expressed in the form of booking limits for each customer class. Inasmuch as this problem deals with demand from different market segments with a different price charged to each segment, one can view determining a room allocation to each segment as a pricing decision.

The hidden price (HP) problem, as we define it here, is characterized by the inability of the hotel’s reservation system to identify the market segment to which a customer belongs at the time that a reservation is made. A simple example of two market segments that are indistinguishable by the booking process would be business people who are traveling to a pre-scheduled meeting and salespeople who are making unscheduled, discretionary calls on customers. In such situations, it is not possible to set booking limits for different customer classes explicitly. However, the room rates that customers are willing to pay is likely to be different in these two market segments. We refer to the maximum rate that a customer is willing to pay for a room as the threshold price. Only by setting a price which excludes one of these market segments can the hotel’s reservation system

distinguish them. Therefore, the pricing policy regulates the sales to different customer classes through the screening of some customers by quoting a price above the price that these customers are willing to pay. The "price elasticity of demand" manifested by the fact that customers from different market segments are willing to pay different amounts for a given room, leads to the issue of setting prices optimally.In this paper, we describe a decision support system for the HP-pricing problem.

In order to differentiate the HP-pricing policy from the overbooking and RP-pricing policies, we briefly review the literature on the two latter issues. The overbooking problem has been extensively researched in isolation from the pricing problem. Rothstein (1971), (1974), (1985) shows the overbooking problem as a stochastic decision process. This means that the decision of how many reservations to take is updated as the rental date draws nearer and actual demand as opposed to forecasted demand manifests itself. A natural question is whether or not such sophisticated analysis is justified by the profitability it creates. Williams (1977) demonstrates the necessity of applying optimization methods to the overbooking and pricing problems as opposed to using simple, approximate decision rules. Ladany (1976) derives models for the overbooking decision process in combination with the RP-pricing decision (also referred to as the "inventory" allocation problem). This latter aspect of the

yield management problem reflects the issue of setting aside a certain amount of hotel capacity for each of several customer classes such as travel groups, single-room customers, double-room customers, etc. Each customer class is given a different room rate. The interest in this problem has persisted in the airline industry for almost thirty years as evidenced by Thompson (1961), Glover et. al. (1982), and many others. See Kimes (1989) for a review of the extensive work in the airline industry on this problem. Lieberman and Yechialli (1978) introduce a discount to a standard room price as a "cost of acquiring a reservation". However, their model does not allow for uncertainty in the response of the market to the discount. Hence, realistic price elasticity of demand that would be characteristic of the HP-pricing policy is not incorporated in any of the aforementioned papers.

While the objective of the overbooking problem is to maximize the occupancy level, the objective of yield management is more generally stated as maximizing profit. The important distinction between these two objectives lies in the fact that the profitability of a hotel is determined by the number of rooms booked as well as by the rates obtained for these rooms and the profit generated by the sale of other products and services to guests. The significant value of products and services beyond the room rental is a distinguishing feature of hotel yield management as opposed to

airline yield management. In the case of the HP-pricing policy, if prices are set too low, the hotel will be filled with low-yield customers, some of whom would have payed a higher price for their rooms. Hence, the capacity of the hotel is not utilized in the most profitable way.

As in the overbooking problem, the pricing policy requires updating in the face of uncertain demand. It is generally more profitable to change prices over time in order to sell some of the hotel capacity at higher prices and to capture a larger number of higher-yield customers. See Orkin (1988) for an overview of these issues. Relihan (1989) describes the nature of this problem in general terms and critiques various proposed approaches to solving it. Belobaba (1989), for the case of RP-pricing policies, and Pfeifer (1990), for the case of HP-pricing policies, give sub-optimal methods for solving these problems by using a concept called "marginal revenue" (MR). The essence of the MR methods is that the marginal value of allocating one unit of capacity (one hotel room or one airline seat) to a higher-yield customer can be based on the difference in threshold prices between customer classes and the probability that the unit of capacity will be demanded by the customer class with the higher threshold price. The MR approach has been applied by Brumelle et. al. (1990) in the case of airline booking when demand is correlated across fare classes. Curry (1990) extends the MR approach to the airline booking problem of in-

corporating the itineraries as well as the fare class of customers in setting booking limits. Both of these papers assume the RP-pricing policies along with rather restrictive assumptions about the demand process. In the current paper, we point out the shortcomings of the MR methods and illustrate the performance of a more accurate model that is derived in Badinelli (1990). The performance of this model in comparison to the MR methods is encouraging in terms of yield as well as computing time. The model and its results apply to HP-pricing policies for hotel bookings, but we expect that the model framework can be extended to the case of airline bookings and RP-pricing policies. While the overbooking and capacity allocation problems have received much research interest in the last twenty years, little work has been done on the problem of dynamic pricing decisions that capitalize on the HP-pricing policy. This suggests our plan of research into yield management which begins with determining the pricing policy, then augmenting this policy with overbooking and capacity allocation considerations.

The pricing problem involves setting the rate charged for a room based on three factors: the date that a customer requests a reservation (booking date), the date of arrival that the reservation calls for (rental date), and the length of stay. There can be other conditions on a rate such as the cancellation policy. Since the focus of the current paper is on the price elasticity of demand and not on the

overbooking problem, we do not include consideration of the cancellation policy. The factors that affect the pricing policy are: the capacity of the hotel, the pattern of demand timing in each market segment (demand process), and the price that customers in each segment are willing to pay (threshold price). The first of these three factors is known exactly. However, the demand pattern can only be forecasted and the pricing policy must be based on this forecast. As time passes and reservations are received, the remaining hotel capacity and the forecast of demand over the time remaining before the rental date can be used to update the pricing decision. The uncertainty under which the decisions must be made and the updating of the decisions over time classifies the problem as a stochastic decision process. As is pointed out by Kimes (1989), since most costs in a hotel operation are fixed, maximizing profit is, for all intents and purposes, equivalent to maximizing revenue and it is the latter objective that we embrace throughout the current paper.

This paper consists of seven sections. In the next section, we identify the essential characteristics of the demand process and the assumptions of the model. In Section 3, we give an overview of the model. The reader who is interested in mathematical details may refer to Badinelli (1990). In Section 4, we compare the model to the traditional MR models. In Section 5, we show an example of the application of

the model and further explain the design of the model. An important aspect of the model is its data and processing requirements. In Section 6 we show that, in application, the user is not burdened by unrealistic data collection and complicated processing. A summary of this research and the avenues for extending the model to more elaborate situations are outlined in the last section.

DEMAND ANALYSIS AND MODEL ASSUMPTIONS

An important and fundamental virtue of the model described herein is the assumption that customers making reservations may identify themselves only by the booking date, the rental date, and the length of stay. Any other information that would reveal the market segment to which a caller belongs and, consequently, suggest a price for quote is not ordinarily communicated in making a reservation. Moreover, any attempt by the hotel industry to secure such information for the purpose of setting prices will inevitably lead to customers routinely misrepresenting themselves in order to obtain the lowest rate possible. We define the reservation process as the time series of reservation calls received by the hotel. We define a demand process as the time series of reservation calls from a given market segment. Since the reservation process is formed by the superposition of the demand processes from all of the market segments, with each market segment reflecting its own threshold price, the hotel's reservation system is confronted with aggregate demand. Due to the differences in the undisclosed threshold prices across the market segments, this aggregate demand manifests random variation in the acceptance/rejection of the quoted price. In spite of this random behavior of aggregate demand, we rely on models of the individual demand processes and estimates of the threshold prices in determin-

ing a pricing policy. The advantage of modeling the demand processes and threshold prices in each market segment derives from the value of such detailed and disaggregated information in constructing a forecast of the reservation process.

A common approach to yield management is to define the service being advertised and the market segment that it is sold to as one in the same. For example, a super-saver fare for a weekend stay focuses on a particular group of vacation travelers who can be identified precisely as the group of people who take advantage of the super-saver offering. This basis for defining market segments leads to the RP-pricing models. The pricing model in the current paper allows for defining market segments in terms of the threshold price, the demand process, the rental date, the length of stay, and the secondary revenue associated with the customers from each segment. For example, business travelers who typically make reservations from one to two weeks prior to the rental date can be modeled as an individual market segment even in the absence of a special advertised rate or package for such customers. While the hotel management may not advertise specific rates for each market segment, we assume that they do reserve the right to adjust the room rate over time based on a forecast of the reservation traffic from the different market segments. This rate-setting decision is precisely the concern of the model described in this paper.

Efforts at capturing demand from different market segments in the airline industry has led to policies that are based on identifying segments with fare classes and controlling the allocation of seats to each class. In a static seat allocation model, the booking limits for each fare class are set in advance of taking reservations and are not updated. Dynamic allocation models update the booking limits as times passes and reservations are received. In an RP-pricing policy with a nested reservation system, seats allocated to lower-fare customers are available to higher-fare customers but not vice versa. In the HP-pricing policy, by adjusting the room rate quoted to customers making reservations, the pricing policy effectively sets booking limits since all segments which have a threshold price below the quoted rate will reject the opportunity to make a reservation. As long as there are unfilled rooms, any request from a customer whose threshold price is at least as high as the quoted rate will be accepted. These properties of the booking policy along with updating the quoted rates over time imbues the policy modeled in the current paper with characteristics of a dynamic allocation model and a nested reservation system.

The model developed in the current paper sets a pricing policy for optimally utilizing the capacity of the hotel for a given rental date. A problem will arise, however, when the length of stay of customers of one rental date causes their occupancy to overlap other rental dates. This makes the

pricing problem a multi-date problem. Hence, while the pricing problem can be solved independently for each rental date, the superposition of the policies for different rental dates may lead to infeasible solutions. We introduce the model in the current paper as a basis for the multi-date model. This extension is the work of future research. However, another point to mention about the multi-date problem is that in many situations, there are critical rental dates such as weekends in resort locations when the hotel is likely to reach full capacity. Bookings for such rental dates will benefit from a yield-management pricing policy. A non-critical rental date is one in which ample capacity is available for all customers. If the critical rental dates are far enough apart in time so that there are few overlapping stays, then the superposition of independently derived pricing policies, using the model described here, should be nearly optimal.

An issue not addressed by the model is the ability to offer different rates to customers from different market segments, as in the RP-pricing policy. A special rate can be given only if a customer identifies himself/herself as being a member of a certain segment such as a convention group, special vacation package tour, etc. Without such identification, if two customers call in a reservation on the same day for the same rental date, the reservation system cannot distinguish them even if their threshold prices are different.

As stated earlier, in such situations, allowing customers to claim to be part of a select group can lead to misrepresentation. In effect, when this happens, the booking policy has caused the marketplace to restructure itself with a higher fraction of low rate-paying customers. In those instances in which a customer can claim a certain rate based on membership in a group, and that membership can be verified, the booking policy should allow for accepting more than one rate. Such a policy would incorporate, as a decision variable, the number of rooms set aside for the special market segment. The price quoted to customers in other segments is a second decision variable. Adding this decision variable to the model is straightforward, but complicates the notation. We leave this extension of the capability of the policy to an enhancement of the model to be presented in a sequel. For the purposes of the current paper, a convention or group-fare market segment is treated like any other segment. Its booking limit is determined when the quoted rate is increased above the rate for this segment.

As the discussion thus far indicates, an important component of the model is the demand forecast. This forecast stems from an analysis of the demand process in each market segment. While information about these market segments generally is not made available through the booking transactions, we assume that the hotel or hotel chain has conducted some kind of market research through surveys as well as through

studies of the patterns of past reservation scenarios and follow-up studies of previous customers. From such studies, the number of market segments can be identified and the demand process for each segment can be estimated. Threshold prices for a segment may be estimated through surveys and experience. However, it is more likely that the most accurate way to estimate threshold prices is by studying the current market rates offered by competitors, since this is how most customers determine their own threshold prices. We assume in the model that each market segment is characterized by a population which makes reservations randomly. For any individual in this population, the decision to make a reservation or not make a reservation and the decision to call in a reservation at a particular time are affected by a variety of factors which are both endemic of the individual's personal affairs as well as the characteristics of the market segment. For example, the market segment consisting of two-day-stay business travelers is formed by the population of all such potential travelers. The decision to book a room at a particular location for any individual in this segment will be influenced by business affairs which crystallize over time until the decision to go or not to go is finalized. Furthermore, the exact date on which the reservation is called in depends, first, on the prerequisite decision to go and, second, on the personal schedule of the individual which may further delay making the reservation call. In effect, the population comprising a given market

segment can be divided into subpopulations based on the time interval in which individuals make their reservations. Such a demand process, composed of a population of individuals who enter the reservation system at random and independently of one another is best described by a poisson process. While the customers in a market segment may call reservations in "at random", the parameters governing this randomness will be characteristics of the market segment. For poisson-distributed demand, this means that the expected demand for each market segment and time interval determines the corresponding demand process. These expected demands are the only parameters of the demand process which must be estimated in order to use the model.

Another important element of the model is the time period prior to the rental date. This time is measured in intervals of varying lengths. The model is based on a discrete time scale implying that each pricing decision will be in effect over an interval of time. Selecting the lengths of the time intervals is influenced by two considerations. First, the desire to update the pricing decisions in order to react to changes in demand rates over time motivates intervals that are short enough to reflect such changes. Second, the amount of data required to estimate the expected demand of a particular market segment and time interval is high when this expectation is very low. Simply stated, the more intervals there are, the more data we need. In some cases, these two

considerations can oppose one another. In cases where sufficient market research has been done, the first consideration given above can determine the lengths of the intervals. The desired interval lengths need not be uniform. For example, when the rental date is months away, typically, the only market segment that is making reservations is the vacation group or convention segment and reservations are being received at a slow rate. Under such conditions, there is no need to update the pricing decision. Hence, the initial price set at the group or convention rate is likely to be in force from six months to three months prior to the rental date. In the three months prior to the rental date, other market segments will show significant activity. During these months a weekly interval may be appropriate until the rental date is a few weeks away. In the last couple of weeks prior to the rental date, some of the highest paying customers will be making their last-minute reservations. Because of the volume of business conducted during this phase and the importance of tracking variations in demand rates, a daily interval may be appropriate. Walk-in business can be modeled as demand during the last time period. This example illustrates a planning horizon of six months or more represented by a time scale having approximately twenty intervals. The number of intervals needed to represent the planning horizon

horizon has a profound impact on the computing time required to determine the optimal pricing policy. The details of this computation are the subject of the next section.

MODEL OVERVIEW

First, we define some notation.

t = rental date
N = total number of time periods in the planning horizon.
k = number of periods in advance of the rental date t that a reservation is made. $k = 1,2,\dots,N$
In other words, we will count time backwards from the rental date.
Recall that the time periods are not necessarily of the same length.
S = total number of market segments.
In the following definitions, unless otherwise stated, i may take on the values $1,2,\dots,S$ and k may take on the values $1,2,\dots,N$.
$L(i)$ = average length of stay for customers in market segment i.
$R(i)$ = the threshold price in \$/day for market segment i.
Without loss of generality, we order the market segments so that $R(1) > R(2) > \dots > R(S)$
$P(i)$ = ancillary profit rate = expected daily profit from a customer from market segment i
in addition to the room charge.

$Q(k,c)$ = room rate quoted in period k to callers requesting reservations when the number of
vacancies = c.
c = remaining unbooked hotel capacity at the beginning of period k.
H = the starting capacity = the number of rooms in the hotel.

First, we note that the room rate quoted to a reservation caller in a given period k, $Q(k,c)$, should be equal to the threshold rate for one of the market segments. To see this, suppose that $Q(k,c)$ is not equal to one of the S threshold prices. That is, suppose $R(i+1) < Q(k,c) < R(i)$ for some market segments i, $i+1$. Then any customers from segments $i+1,\dots,S$ will reject the offered rate and customers from segments $1,2,\dots,i$ will accept the rate. Clearly, if we wish to capture demand from segments $1,2,\dots,i$, we can do so more profitably by setting $Q(k,c) = R(i)$. We define $i(k,c)$ as the rate class of the rate chosen for $Q(k,c)$; that is, $Q(k,c) = R(i(k,c))$. In effect, by identifying $i(k,c)$ for each possible value of c we set nested booking limits for each market segment. As more rooms are sold, c decreases, motivating an increase in $Q(k,c)$. A particular value of c at which $Q(k,c)$ increases represents a booking limit for customers with lower threshold prices.

For each reservation taken from market segments 1, 2,...,i(k,c), the expected yield is the number of days stay times the revenue per day or, L(j) (R(i) + P(j)), j=1,...,i. This expression shows the actual yield that is gained by the pricing decision in period k for a given customer demand from segment j. However, the pricing decision must be made prior to these demands being known. Therefore, the objective of the decision is to maximize the EXPECTED yield in period k, which can be computed from the probability distributions of demand. We rely on the demand forecast to provide the probability distribution of demand in each market segment. We can state the optimization problem that represents the problem of setting an optimal pricing policy over the entire planning horizon as follows: Starting with c = H available rooms in period N, set the quote price for each reservation request depending on the time of the request and the hotel's situation at that time. In effect, we must determine the optimal values for Q(k,c) for k=1,...,N and for c=1,..., H. In other words, we wish to maximize the total expected yield over all N time periods by manipulating the quoted prices over these intervals subject to the constraints that force the number of rooms sold in any period to be no greater than the remaining number of unbooked rooms in the hotel. The sequential nature of the decision process that is represented by this problem suggests that a dynamic programming approach to solving the problem is most efficient.

A dynamic programming reformulation follows naturally. In order to apply the dynamic programming approach, we start with period 1 with one vacant room (c = 1) and find the optimal quote price. We then proceed backwards in time to period N and upwards in capacity to H. At each iteration we build on the policy obtained thus far by computing the optimal quote price for one more vacant room. This optimal rate is determined as the rate which maximizes the total expected yield from the current period until the rental date. In the same way, the decision rules obtained in period k will be used to build the solutions for period k+1, etc.

COMPARISON TO MR MODELS

The form of the HP-pricing problem appears to be more complex than the formulation found in MR models such as that of Pfeifer (1989). In fact, the MR models are simpler. However, embedded in this relative simplicity of form is an approximation that, as we will see later, causes a serious degradation in performance. We can illustrate the nature of this approximation with a simple example. Consider a situation in which the remaining capacity is four rooms and there are two market segments having threshold prices of $60 and $40. Assume that the next caller will be someone from the $60 segment and that after this caller there will be two more from the $60 segment and three more from the $40 seg-

ment. In keeping with the assumed structure of the pricing problem that is characteristic of MR methods, we assume that once the quote price is raised, it cannot be lowered. The question to be answered then, is when to raise the quote price to $60. This can be done at the time of any one of the six incoming calls. A seventh option would be to not raise the price at all. Using the MR approach of Pfeifer (1989), the pricing policy can be determined without any further information about demand. However, such an approach is overly simplistic. In fact, the scenario of demand from each segment, and not just the total demand from each segment, influences this decision. In order to see why this is so, consider the pattern of incoming reservation calls from the two segments shown in Figure 1.

$60 segment	x				x	x
$40 segment		x	x	x		
Sequence #	1	2	3	4	5	6

c=4

Time of price increase to $60	Yield
1	$180
2	$160
3	$200 *
4	$180
5	$160
6	$160
no increase	$160

Figure 1 - Demand Pattern A

As one can verify by computing the yield from each of the seven options, the optimal decision is to raise the price from $40 to $60 with the third caller. Now suppose the demand pattern is as shown in Figure 2. Note that the next caller is still from the $60 segment and that the total number of callers from each segment is the same as before. However, the scenario of demand after the first caller is different. Once again, we can determine the optimal time to change the price to $60 by computing the yield for each of the seven options. In this case, the optimal decision is to raise the price with the first caller.

$60 segment	x	x	x			
$40 segment				x	x	x
Sequence #	1	2	3	4	5	6

c=4

Time of price increase to $60	Yield
1	$180 *
2	$160
3	$140

4	$120
5	$160
6	$160
no increase	$160

Figure 2 - Demand Pattern B

Now suppose that demand is not completely predictable, but can only be forecasted, with the result that the probability distribution of demand scenarios can be specified. Suppose

the demand forecast states that the probabilities of pattern A shown in Figure 1 and pattern B shown in Figure 2 are .5 and .5 respectively. Computing the expected yield for each of the seven pricing options determines the optimal pricing policy as that of raising the price with the first caller. A different forecast will generally result in a different pricing policy. Suppose that the probabilities of patterns A and B are .8 and .2, respectively. In this case, the optimal time to raise the price is with the third customer. In order to preclude this dependence of the pricing policy on the forecast of demand scenarios as opposed to just the simpler forecast of total demand from each segment, it is common to see assumptions in the research literature about the allowed patterns of demand. Typically, one assumes that the lower fare customers make their reservation calls before the higher fare customers. The lack of detailed studies on actual demand scenarios has kept this assertion unverifiable although recent studies cited by Lee (1990) seem to cast doubt on the assumption’s validity. A confounding influence in any attempt to measure demand patterns is that the pricing policies extant act to influence those demand patterns. For example, the offer of discounted rates to callers who make their reservation more than sixty days in advance will make the occurrence of lower-fare demand prior to higher fare demand a self-fulfilling prophecy. In the model developed by Badinelli (1990), there is no assumption regarding

the timing of lower fare demand relative to higher fare demand.

The astute reader may realize from the examples given above that the optimal pricing policy may involve raising prices, then lowering prices later, depending on the comparison of actual demand relative to the forecast. In MR models, such a reversal is usually disallowed. One might think that by repetitively applying a MR model with updated forecasts that an optimal sequence of decisions would result. However, this is not true. For example, if a pricing policy raised the quote price six weeks prior to the rental date in anticipation of a significant amount of higher fare demand, this decision may be reversed later on if this higher fare demand fails to materialize as originally expected. In effect, the future price reduction is contingent on the state of the system at a point in time in the future, where that state can be expressed as the amount of demand that has already taken place or, equivalently, as the number of vacant rooms remaining at that time. Carrying this possibility further, we can see that the wisdom of raising the quote price six weeks prior to the booking date depends on the possibility of changing prices later. In order to achieve an optimal sequence of pricing decisions, one must not only make use of updated forecasts as time proceeds, but also, in making a pricing decision at one point in time, the possible effects of future pricing decisions must be incorporated in the cur-

rent decision. This means that optimal pricing decisions cannot be made "as you go along", but that an optimal pricing policy must specify all future pricing decisions, contingent on actual demand performance. That is, just as the optimal pricing decision can be determined only by knowing the forecast of demand scenarios as opposed to forecasts of total demand, the optimal pricing decision at any point in time can only be determined by conditioning the decision on future pricing decisions. In effect, the optimal pricing policy must be specified as a sequence of state-dependent pricing decisions DETERMINED JOINTLY.

The point of this exercise is to show that unless rather restrictive conditions are assumed for the demand scenarios, the pricing decision must be based on a forecast of not only the total demand from each segment, but the pattern of that demand. Also, the optimal decision at any point in time must be conditioned on future pricing decisions that might be made. The effect of these considerations in setting the pricing policy necessitates the use of analysis more complicated that the MR approach. The name for this analysis is dynamic programming. Resistance to the use of dynamic programming has been based on the misconception that it is a) too costly in terms of computing time and b) too complex to explain to managers. In this paper, we assert that the dynamic programming formulation of the hotel HP-pricing problem as formulated in Badinelli (1990) actually requires less

computing time than the MR approach and results in significantly higher yield. Furthermore, we assert that the model can be implemented by hotel managers without an understanding of dynamic programming beyond the explanation of the nature of the optimal pricing policy given above.

COMPUTATION OF THE OPTIMAL POLICY

The HP-pricing policy is computed by starting with period 1, the last period before the rental date, and proceeding backward in time to period N. Each iteration involves a pricing decision which is conditioned on the remaining capacity. The optimal quoted rate for period k, under the condition that c empty rooms remain is $Q(k,c) = R(i(k,c))$. This decision must be computed for all possible capacities (between zero and the number of rooms in the hotel). This optimal solution for Q(k,c) is based on the optimal sequence of pricing decisions after period k. At iteration k, this optimal sequence of decisions after period k would have been computed in previous iterations. In effect, then, the dynamic programming solution results in the computation of DECISION RULES which prescribe the optimal quoted rate for each time period, conditioned on the remaining capacity that exists at the beginning of each time period. In this way an optimal sequence of decisions from period N through period 1 can be constructed.

Table 1 shows an example of an optimal solution. The situation depicted in this example is that of a 40-room hotel. There are three market segments and five periods. The small dimensions and simplistic parameter values of this example are chosen for the sake of exposition. Table 1 shows the threshold prices and the demand rates for each segment and time period. The solution is stated in the form of the decision rules which identify, for each of the five time periods, an optimal quoted rate for every possible value of remaining capacity.

segment	expected demand rates 1	2	3	4	5	threshold rates	avrg stay	ancillary profit
1	10	5	0	0	0	70	1	0
2	0	5	10	5	0	60	1	0
3	10	10	10	10	10	50	1	0

Period	Remaining Capacity	Quoted Rate(optimal)
5	30-40	50
5	1-29	70
4	33-40	50
4	12-32	60
4	1-11	70
3	32-40	50
3	12-31	60
3	1-11	70
2	20-40	50
2	13-19	60
2	1-12	70
1	11-40	50
1	1-10	70

Table 1: Decision Rules - Sample Problem

One can see from Table 1, that by determining Q(k,c) for all k and c, we have set booking limits for each of the three segments. For example, in period 1, the booking limit for segments 2 and 3 is 30 rooms. The booking limit for segment

1 is 40 rooms. By the nature of the HP-pricing problem, these booking limits are necessarily nested. Table 1 also shows that in periods 3,4, and 5, there are situations in which we would not accept any reservations, even though we have vacancies. In these periods, when $Q(k,c) = \$70$, we are effectively holding rooms for high-yield demand in later periods. It is also worth noting that this optimal solution prescribes lowering prices in certain situations. For example, if there are 25 vacant rooms in period 3, the reservation system would quote a price of $60 to any caller. If by period 2 there are still 25 vacant rooms, the quote price is reduced to $50. One can see that such a policy makes sense given the forecast of demand shown in Table 1. In an MR model such as that of Pfeifer (1989) such a price decrease is not possible.

A by-product of the Badinelli model is the computation of the expected yield that results from using the optimal policy. For the example shown in Table 1, the expected yield is is $2345. By comparison, the same problem solved using an MR model results in a policy with an expected yield of $2274.

IMPLEMENTATION ISSUES

Techniques such as dynamic programming often have been criticized for their complexity, massive data requirements, long computing time requirements, and unintelligibility to most practitioners. See Kimes (1989) for such critiques of dynamic programming as applied to the yield management problem. In light of these criticisms, it is important that we delineate the role played by the hotel management in making use of the model presented in the current paper. First, let us examine the data requirements. A basic requirement for applying the model is the identification of the market segments. Every hotel can identify at least three segments: vacation travelers, group/convention travelers, and business travelers. Some hotel chains have already identified many more segments through their market research. One large chain has identified approximately twenty market segments. The model is flexible in that it will provide as much detail in the solution as the hotel management has identified in the market. Therefore, many hotels are in a position to provide this basic information to the model. Given an identification of the market segments, the inputs to the model that must be provided by the hotel management are:

1. threshold rates for each market segment: R(1), R(2), ..., R(S)
2. ancillary profit rate for each market segment: P(1), P(2), ..., P(S)
3. average length of stay for each market segment: L(1), L(2), ..., L(S)
4. expected demand rates for each market segment and

time period

As described earlier in the paper, the threshold rates can be estimated either through market surveys and experience or by examining the competitive situation. The ancillary profit rates can be estimated from the hotel's own historical data on meal, gift, room service, and other service charges as a function of market segment. The length of stay is important to the yield. All other factors being equal, the length of stay can define a market segment. For example, business travelers who stay for two days can be identified as a segment separate from business travelers who stay for three days. However, if both of these market segments have the same threshold rate, there would be no benefit to doing this. The model will arrive at the same pricing policy and the same total yield if these two segments are combined and the length of stay for the combined segment is set equal to the average length of stay for all customers in this group. Therefore, the estimates of average length of stay do not

necessarily require extensive data analysis. Finally, the probability distributions of demand that are used in solving the HP-pricing problem are needed. Since these distributions are poisson, only the mean demand rate for each market segment and time period is needed. That is, the poisson probability distribution depends on only one parameter, its mean. All of the properties and computed probabilities of this distribution can be obtained by the computer program once the mean is known. Therefore, the forecasting input to the model consists of the expected demand for each market segment in each time period listed above. As stated earlier, historical data from the hotel's reservation system can be used to estimate these expected values. It is important to note that the numbers being estimated represent DEMAND, not SALES. The expected demand rate is the average number of reservation calls that are received from a particular market segment in a given period. Due to the price being quoted or the possibility of the hotel being fully booked at any given time, not all reservation calls result in sales. Therefore, the data used to estimate the mean demand rates must come from the reservation traffic, and not from the sales figures.

The data requirements given above are within the abilities of even a small hotel to provide. It must be emphasized that all other aspects of the computation required to solve the HP-pricing problem are transparent to the hotel manager.

The problem would be solved once for each rental date of concern to the hotel manager. The computation time required is modest, even to the point that a small hotel can apply the model using a typical personal computer. To prove this point, the problem solution shown in Table 1 was computed via a program written in FORTRAN and run on a Zenith-286 PC.

The time to generate the solution was 7.09 seconds. The computation time for the problem increases with the number of segments, S, the number of time periods, N, and the number of rooms in the hotel, H. Typical values for both S and N are, perhaps, 10. A typical value for H for a small hotel that would have access to only a PC would be 100. Problems with these parameters were run on the Zenith-286 PC. The average computation time was 30 seconds. This, we claim, is not excessively long, especially in light of the fact that the computation can be run at any time of the day or night, whenever it is most convenient to use the hotel's PC.

SUMMARY AND EXTENSIONS

We have shown in this paper a model that possesses reasonably good potential for application that will enhance hotel yield management. This model also represents a framework for extensions that will solve more elaborate yield management problems. Among these extensions are the introduction of multiple room classes and allocation decisions, overbooking considerations, and the overlapping of stays. With the

introduction of multiple room classes and overbooking considerations the problem is still amenable to a dynamic programming approach. However, the number of decision variables is increased and the the data requirements for demand forecasting are more extensive. These extended problems should still be within the abilities of mid- to large-size hotels or chains. The introduction of overlapped stays implies that the pricing policy for one rental date is intertwined with that of other rental dates. Solving this problem will require constructing a more elaborate optimization problem. In order to make a solution more tractable, computationally, we expect that subproblems in addition to the dynamic programming formulation will be needed to bound or approximate solutions. Work on this problem is under way. Finally, one can see an enhancement to the model that is necessitated by the lack of demand data in some situations. In the case of a new hotel or in situations of a rapidly changing market and competitive environment, historical demand data may not produce good estimates of future demand rates. In such cases, as the planning horizon unfolds, the hotel manager acquires fresh demand data which can be incorporated in updated estimates of demand rates. An optimal policy for such a situation would be considered an "adaptive control policy." The formulation would, once again, be a dynamic programming one.

In summary, we have presented a seminal formulation of the hotel yield management problem that is more robust and profitable than the MR models. We have shown that this model is computationally efficient and requires data that is obtainable by most hotel managers. The dynamic programming formulation which incorporates a set of market segments, each characterized by a threshold price, ancillary profit, and average length of stay can be extended to include more elaborate conditions on the pricing, overbooking, and capacity allocation decisions.

References

Glover, F., Glover R., Lorenzo, J., McMillan, C. (1982), The passenger mix problem in the scheduled airlines. INTERFACES 12, 73-79.

Badinelli, R. (1990), A Time-Based Decision-Support Model for Hotel Yield Management.Working Paper, Virginia Polytechnic Institute & State University.

Belobaba, P., (1989), Application of a Probabilistic Decision Model to Airline Seat Inventory Control.OPERATIONS RESEARCH 37, 183-197.

Brumelle, S., McGill, J., Oum, T., Sawaki, K., Tretheway, M. (1990), Allocation of airline seats between stochastically dependent demands. TRANSPORTATION SCIENCE 24, 183-192.

Curry, R. (1990), Optimal airline seat allocation with fare classes nested by origins and destinations. TRANSPORTATION SCIENCE 24, 193-204.

Glover, F., Glover R., Lorenzo, J., McMillan, C. (1982), The passenger mix problem in the scheduled airlines. INTERFACES 12, 73-79.

Kimes, S. (1989), Yield management concepts and solution techniques: application to the lodging industry. Presented at the national meeting of the Decision Sciences Institute, November, 1989.

Ladany, S. (1976), Dynamic operating rules for motel reservations. DECISION SCIENCES 7, 829-840.

Lee, A. (1990), Reservations forecasting for yield management. Paper presented at ORSA/TIMS Meeting in Philadelphia, October.

Lieberman, V., Yechialli, U. (1978), On the hotel overbooking problem: an inventory problem with stochastic cancellations. MANAGEMENT SCIENCE 24, 1117-1126.

Orkin, E. (1988), Boosting your bottom line with yield management. CORNELL HOTEL AND RESTAURANT ADMINISTRATION QUARTERLY February, 52-56.

Pfeifer, P. (1989), The airline discount fare allocation problem. DECISION SCIENCES 20, 149-157.

Relihan, W. (1989), The yield-management approach to hotel-room pricing. CORNELL HOTEL AND RESTAURANT ADMINISTRATION

QUARTERLY 30, 39-45.

Rothstein, M. (1971), An airline overbooking model. TRANSPORTATION SCIENCE 5, 180-192.

Rothstein, M. (1974), Hotel overbooking as a markovian sequential decision process. DECISION SCIENCES, 5, 389-394.

Rothstein, M. (1985), Operations research and the airline overbooking problem. OPERATIONS RESEARCH 33, 237-248.

Thompson, H. (1961), Statistical problems in airline reservation control. OPERATIONAL RESEARCH QUARTERLY 12, 167-185.

Williams, F. (1977), Decision theory and the innkeeper: an approach for setting hotel reservation policy. INTERFACES 7, 18-30.

- - - - - - - - - - - - -

EDITORIAL

Welcome to this first issue of JIAHR, "The Journal of the International Academy of Hospitality Research." The Fellows of the Academy, the editors and the publisher look forward to serving you.

JIAHR, we believe, is the first journal dedicated to serving the worldwide community of men and women who conduct basic and applied research in the field of hospitality and tourism.

Our idea, quite simply, is to communicate their reports on the results of their research. This first issue carries just such a report in Ralph Badinelli's and Michael Olsen's outstanding paper "Hotel Yield Management Using Optimal Decision Rules." Future issues may carry occasional reviews, commentary, and other information of importance to those in hospitality research, but the basic issue each time will be a single research paper.

It may be that JIAHR also will serve a larger community of persons in the hospitality industry -- undergraduate students, teaching faculty, professionals in the hospitality and tourism industry who are interested in the research being reported. If so, that will be satisfying to us all.

JIAHR is also one of the first scholarly journals sent to its subscribers electronically -- and, in fact, apparently is the first to be sent on line to paid subscribers. This is important, but we hope not too important. The technical aspects of an electronic journal are so interesting, absorbing, and sometimes frustrating that it is easy to let them take precedence over more important scholarly concerns, especially in the beginning phases of the journal. May this not happen to JIAHR!

An electronic journal has both advantages and disadvantages. It is certainly much less costly to produce and distribute the journal electronically than on paper, and that's an enormous advantage. Because it is electronic, JIAHR can, and will, publish its papers instantly when they are judged ready for publication; there will be no wait for a quarterly mail date, or a press run, or schedules of airplanes, ships and trucks to bring it to you. Later, we plan to archive these papers electronically in such a way that you can access them from wherever you are, and search them, and use them easily and conveniently in your daily work. We also plan to encourage reader feedback to articles -- instantly in electronic messages, in a way that print journals can't hope to accomplish. That can lead to electronic "conferencing" via e-mail on hospitality issues. These are advantages that would be impossible without the electronic medium.

As a pioneer in electronic publication, however, JIAHR will experience -- and already is experiencing -- problems not found with printed publications: problems with e-mail addresses, graphics (we can publish nothing more complicated than tables at present), standards, internal distribution within universities and libraries, subscribers unfamiliar with the technology, and many more. Many in hospitality research are not yet comfortably equipped to receive JIAHR. We are working with such individuals, cajoling and encouraging them to obtain e-mail addresses, to experience the enormous benefits of this new world of communication not only by receiving this journal but by establishing instant communication with colleagues the world over. We are working with libraries as they establish their own procedures to make JIAHR (and other electronic journals which will follow) available to their patrons. By encouraging these people to accommodate themselves to JIAHR, we think we are helping to lead the hospitality industry -- and all of society -- into an exciting new world of communication, a new age of information. Hence the problems we face are opportunities.

In the meantime -- to work with reality -- until they are satisfactorily "hooked up" electronically, subscribers who cannot yet receive the journal by BITNET or the Internet will be sent it quarterly on floppy disks.

As individuals and libraries gear up mentally and technologically to participate in this new age, we take pride in playing a trail-blazing role. We hope that all subscribers -- individual and institutional -- will play this role, too. As a subscriber, you are asked to send messages to the editor (JIAHRED@VTVM1.Bitnet and JIAHRED@VTVM1.CC.VT.EDU on the Internet) telling of your experience with the journal: what you like, what you dislike, problems, solutions, questions, thoughts and considerations about the hospitality issues raised in JIAHR. Questions about subscriptions may be sent to SCHOLAR@VTVM1 (SCHOLAR@VTVM1.CC.VT.EDU on the Internet). If you're not yet comfortable communicating electronically, send us a letter (JIAHR, Scholarly Communications Project, Virginia Tech, 1700 Pratt Drive, Blacksburg, VA 24061-0506).

Authors are encouraged to submit papers (instructions to authors are found elsewhere in this issue) for consideration. Electronic communication allows quick review and feedback that print journals would have trouble matching.

Together, authors, editors, the publishing staff and above all the readers can work out the technical problems while we improve the journal's service and content. In doing so, we will help build knowledge in the hospitality field while developing electronic communication to benefit not only ourselves but those in other, less favored disciplines that will follow our lead.

- - - - - - - - - - - -

Price of Annual Subscription: Institutional $30; Individual $20; Student (with professor's attestation) $10. Send to

Scholarly Communications Project, Virginia Tech, 1700 Pratt Drive, Blacksburg, Va. 24061-0560.

This journal is registered with the Copyright Clearance Center, Inc., 27 Congress Street, Salem, MA 01970, U.S.A. Limited duplication is permitted for academic or research purposes only. Published for the International Academy of Hospitality Research by the Scholarly Communications Project of Virginia Tech.

It is our plan that JIAHR subscribers will be able to obtain abstracts and/or full texts of papers previously published by the journal by electronic means. This service will be established after several issues have been published.

Research scholars are invited to submit basic and/or applied research papers to JIAHR covering all aspects of hospitality and tourism. Articles should not exceed 10,000 words and should include an abstract, key words and references. Papers will be refereed by at least two persons. Electronic communication assures prompt review and decision on acceptance. Published papers may include tables (but not graphs or art work) and should be submitted single spaced as electronic files in

ASCII format to JIAHRED@VTVM1.Bitnet or JIAHRED@VTVM1.CC.VT.EDU on the Internet. Authors without access to Bitnet or the Internet may submit papers on disks in ASCII format. Papers should be written according to the Publication Manual of the American Psychological Association. For further information and/or guidelines about paper submissions, contact Managing Editor, JIAHR, 18 Hillcrest Hall, Blacksburg, VA 24061-0429, USA; telephone 703/231-5515 or 703/231-4922, USA; or send e-mail to Bitnet or Internet addresses above.