Yield Management

Serguei Netessine[1]
The Wharton School
University of Pennsylvania

Robert Shumsky[2]
W. E. Simon Graduate School of Business Administration
University of Rochester

February 1999, revised February 2002.

[1] netessin@wharton.upenn.edu
[2] shumsky@simon.rochester.edu

1. Introduction

A variety of concepts and analytical tools fall under the label *yield management*. The term is used in many service industries to describe techniques to allocate limited resources, such as airplane seats or hotel rooms, among a variety of customers, such as business or leisure travelers. Since these techniques are used by firms with extremely perishable goods, or by firms with services that cannot be stored at all, these concepts and tools are often called *perishable asset revenue management* or simply *revenue management*.

The techniques of yield management are relatively new – the first research to deal directly with these issues appeared less than 20 years ago. These days, yield management, including overbooking and dynamic pricing, has been an enormously important innovation in the service industries. American Airlines credits yield management techniques for a revenue increase of $500 million/year and Delta Airlines uses similar systems to generate additional revenues of $300 million per year.[3] While the airlines are the oldest and most sophisticated users of yield management, these practices have begun to appear in other service industries. For example, Marriott hotels credits their yield management system for additional revenues of $100 million per year. All of these revenue increases have been achieved with relatively small increases in capacity and costs.

This note describes yield management's basic concepts and provides details about a particular application. Our real-world example is provided by the Hyatt Regency hotel at 125 East Main Street in Rochester, NY. The hotel has 210 King/Queen rooms used by both business and leisure travelers. The hotel must decide whether to sell rooms well in advance at a relatively low price (i.e. to leisure travelers), or to 'hold out' and wait for a sale at a higher price to late-booking business travelers. Before we examine a specific solution to this problem, we will first identify the environments in which yield management techniques are most successful.

2. Where and Why Firms Practice Yield Management

Business environments with the following five characteristics are appropriate for the practice of yield management (in parentheses we apply each characteristic to the Hyatt hotel):

[3] Andrew Boyd, "Airline Alliance Revenue Management". OR/MS Today, Vol. 25, October 1998.

1. It is expensive or impossible to store excess resource (we cannot store tonight's room for use by tomorrow night's customer).
2. Commitments need to be made when future demand is uncertain (we must set aside rooms for business customers – "protect" them from low-priced leisure travelers - before we know how many business customers will arrive).
3. The firm can discriminate among customer segments, and each segment has different demand curves (purchase restrictions and refundability requirements help to segment the market between leisure and business customers. The latter are more indifferent to the price.).
4. The same unit of capacity can be used to deliver many different products or services (rooms are essentially the same, whether used by business or leisure travelers).
5. Producers are profit-oriented and have broad freedom of action (in the hotel industry, withholding rooms from current customers for future profit is not illegal or morally irresponsible. On the other hand, such practices would be questionable in emergency wards or with organ transplants).

Given these characteristics, how does Yield Management work? Suppose that our hotel has established two fare classes or buckets: *full price* and *discount price*. The hotel has 210 rooms available for March 29 (let us assume that this March 29 is a Monday night). It is now the end of February, and the hotel is beginning to take reservations for that night. The hotel could sell out all 210 rooms to leisure travelers at the discount price, but it also knows that an increasing number of business customers will request rooms as March 29 approaches and that these business customers are willing to pay full price. To simplify our problem, let us assume that leisure demand occurs first and then business demand occurs. Hence we must decide how many rooms we are willing to sell at the leisure fare or in other words, how many rooms shall we protect (i.e., reserve) for the full price payers. If too many rooms are protected, then there may be empty rooms when March 29 arrives. If too few are protected, then the hotel forgoes the extra revenue it may have received from business customers.

Notice we assume that the hotel can charge two different prices for the same product (a room). To separate the two customer segments with two different demand curves, the hotel introduces *price discrimination.* In the hotel example, we assume that the firm can charge different prices to business and leisure customers[4]. In order to differentiate between these two groups, a firm often introduces booking rules. For example, a

Saturday-night stay may be required to receive a discounted room on Monday, since business travelers are less likely to stay over the weekend. Again, leisure customers are more price-sensitive so the hotel wishes to sell as many rooms to business customers at a higher price as possible while keeping room utilization high.

3. Booking Limits and Protection Levels

Before we look at the mathematics that will help us to make this decision, some vocabulary would be helpful. We define a *booking limit* to be the maximum number of rooms that may be sold at the discount price. As we noted previously, we assume that leisure customers arrive *before* business customers so the booking limit constraints the number of rooms that these customers get: once the booking limit is reached, all future customers will be offered the full price. The *protection level* is the number of rooms we will *not* sell to leisure customers because of the possibility that business customers might book later in time. Since there are 210 rooms available in the hotel, and just two fare classes, in our example,

booking limit = 210 – the protection level.

Therefore, the hotel's task is to determine either a booking limit or a protection level, since knowing one allows you to calculate the other. We shall evaluate the protection level. Suppose the hotel considers protection level '*Q*' instead of current protection level *Q+1* (*Q* might be anything from 0 to 209). Further, suppose that 210-*Q-1* rooms have already been sold (see Figure 1). Now a prospective customer calls and wishes to reserve the first 'protected' room at the discount price.

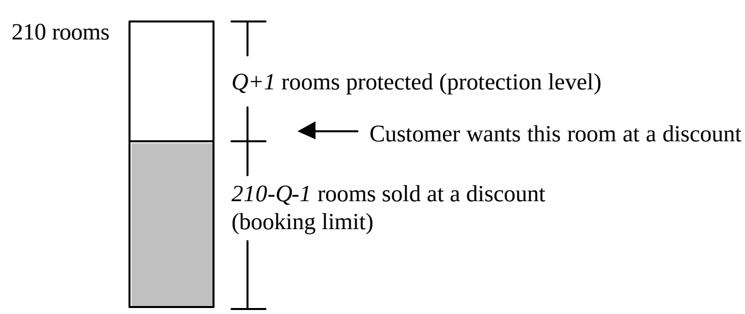

Figure 1. Protection Level and Booking Limit in the Hotel

---

[4] Finding the appropriate price for each customer segment is not in the scope of this note. We assume that prices are established in advance.

Should the hotel lower the protection level from $Q+1$ to Q and therefore allow the booking of the (*Q+1)*th room at the discount price? Or should it refuse the booking to gamble that it will be able to sell the very same room to a full price customer in the future? The answer, of course, depends on (i) the relative sizes of the full and discount prices and (ii) the anticipated demand for full price rooms. The decision is illustrated as a decision tree in Figure 2.

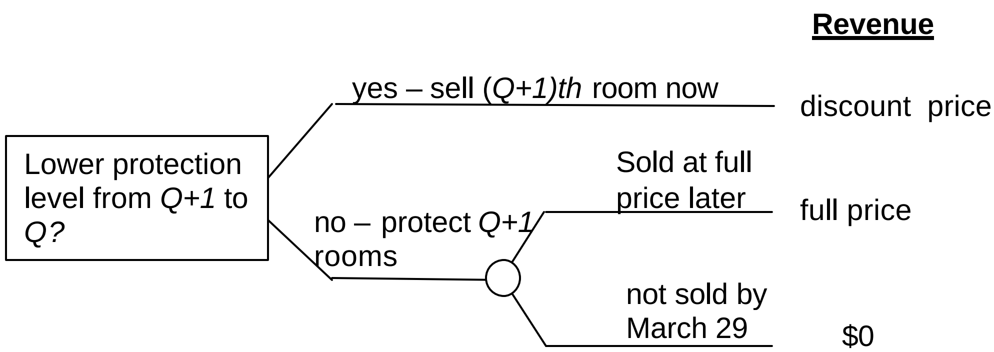

Figure 2. Deciding the protection level

In the next section we associate numbers with this decision and find the optimal protection level, Q^*.

4. Solving the Problem

To determine the value of each branch of the decision tree in Figure 2, we need to know the probability for each 'chance' branch and the values at the end of the branches. Suppose that the discount price is $105 per night while the full price is $159 per night. To find the probability on each branch, define random variable D to represent the anticipated demand for rooms *at the full price*. The hotel may estimate the distribution of D from historical demand, as well as from forecasts based on the day of the week, whether there is a holiday, and other predictable events. Here we will assume that the distribution is derived directly from 123 days of historical demand, as shown in Table 1 below.

In the table, the 'Cumulative Probability' is the fraction of days with demand at or below the number of rooms in the first column (Q).

Demand for rooms at full fare (Q)	# Days with Demand	Cumulative Probability *F(Q)* *=Prob{D≤Q}*
0 - 70	12	0.098
71	3	0.122
72	3	0.146
73	2	0.163
74	0	0.163
75	4	0.195
76	4	0.228
77	5	0.268
78	2	0.285
79	7	0.341
80	4	0.374
81	10	0.455
82	13	0.561
83	12	0.659
84	4	0.691
85	9	0.764
86	10	0.846
above 86	19	1.000
Total	123	1.000

Table 1. Historical demand for rooms at the full fare.

Now consider the decision displayed in Figure 2. If we decide to protect the (*Q+1)*th room from sale, then that room may, or may not, be sold later. It will be sold only if demand *D* at full fare is greater than or equal to *Q+1*, and this event has probability *1-F(Q).* Likewise, the protected room will not be sold if demand is less than or equal to *Q*, with probability *F(Q).* Figure 3 shows our decision with these values included.

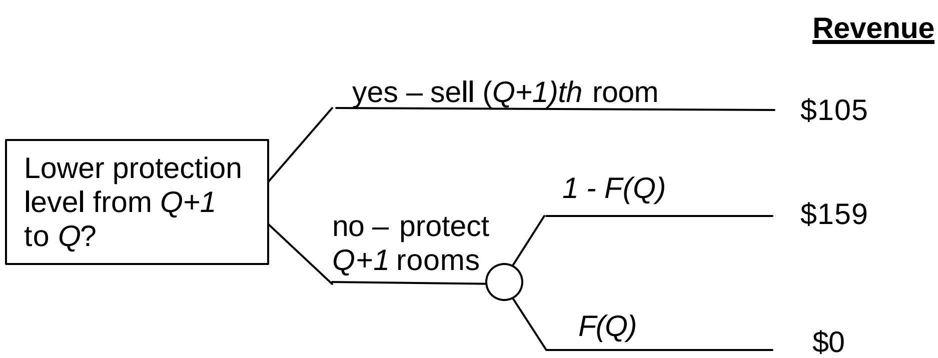

Figure 3. Booking Limit Decision with Data

Given Figure 3, we can calculate the value of lowering the protection level from $Q+1$ to Q. Lowering the protection level results in selling the $(Q+1)th$ room at a discount which guarantees revenue of $105. Protecting $Q+1$ rooms has an expected value equal to:

$$(1 - F(Q))(\$159) + F(Q)(\$0) = (1 - F(Q))(\$159)$$

Therefore, we should lower the booking limit to Q as long as:

$$(1 - F(Q))(\$159) \leq \$105$$

or

$$F(Q) \geq (\$159 - \$105) / \$159 = 0.339.$$

Now, $F(Q)$ is the third column in Table 1. We simply scan from the top of the table towards the bottom until we find the smallest Q with a cumulative value greater than or equal to 0.339. The answer here is that the optimal protection level is Q^*=79 with a cumulative value of 0.341. We can now evaluate our booking limit: 210 - 79 = 131. If we choose a larger Q^*, then we would be protecting too many rooms thereby leaving too many rooms unsold on average. If we set Q^* at a smaller value, we are likely to sell too many rooms at a discount thereby turning away too many business customers on average.

5. A General Formula

The solution described above is an example of a standard technique that was developed for the airline industry. The technique was named “Expected Marginal Seat Revenue” (EMSR) analysis by Peter Belobaba at MIT[5]. In our example we had two fare classes with prices r_L and r_H (r_H is the higher price, $159, while r_L is the lower price, $105). We also had a random variable, D, to represent the distribution of demand at the high fare. Since there are just two fare classes, the optimal booking limit for low fare class is equal to the total capacity minus Q^*.

The EMSR analysis can actually be described as a newsvendor problem. We make a fixed decision, Q, the protection level, then a random event occurs, the number of people requesting a room at the full fare. If we protect too many rooms, i.e., $D<Q$, then we end up earning nothing on Q-D rooms. So the overage penalty is r_L per unsold room. If we protect too few rooms, i.e., $D>Q$, then we forgo r_H-r_L in potential revenue on each of the

---

5 Peter P. Belobaba "Application of a Probabilistic Decision Model to Airline Seat Inventory Control". Operations Research Vol. 37 No. 2, 1989.

D-Q rooms that could have been sold at the higher fare so the underage penalty is r_H-r_L. The critical ratio is

$$\frac{B}{B+C} = \frac{r_H - r_L}{r_L + r_H - r_L} = \frac{r_H - r_L}{r_H}.$$

From the newsvendor analysis, the optimal protection level is the smallest value Q^* such that

$$F(Q^*) \geq \frac{B}{B+C} = \frac{r_H - r_L}{r_H}.$$

6. Overbooking

Another important component in the yield management toolbox is the use of *overbooking* when there is a chance that a customer may not appear. For example, it is possible for a customer to book a ticket on an airline flight and not show up for the departure. If that is the case, the airline may end up flying an empty seat resulting in lost revenue for the company. In order to account for such no-shows, airlines routinely overbook their flights: based on the historical rate of no-shows the firm books more customers than available seats. If, by chance, an unusually large proportion of the customers show up, then the firm will be forced to 'bump' some customers to another flight. Hotels and rental car agencies also overbook. When determining the optimal level of overbooking, the calculation is similar to the calculation used for yield management. The optimal overbooking level balances (i) lost revenue due to empty seats and (ii) penalties (financial compensation to bumped customers) and loss of customer goodwill when the firm is faced with more demand than available capacity.

Let X be the number of no-shows and suppose we forecast the distribution of no-shows with distribution function $F(x)$. Let Y be the number of seats we will overbook, i.e., if the airplane has S seats then we will sell up to $S+Y$ tickets. As in the newsvendor model, define the underage penalty by B and the overage penalty by C. In this case C represents ill-will and the net penalties that are associated with refusing a seat to a passenger holding a confirmed reservation (the ‘net’ refers to the fact that the airline may collect and hold onto some revenue from this passenger as well). Here, B represents the opportunity cost of flying an empty seat. To explain further, if $X>Y$ then we could have sold X-Y more seats and those passengers would have seats on the plane. So B equals the price of a ticket. If $X<Y$ then we need to bump Y-X customers and each bump has a net cost of C.

Thus, the formula for an optimal number of overbooked seats takes the following familiar form: the optimal number of overbooked seats is the smallest value *Y** such that

$$F(Y^*) \geq \frac{B}{B+C}.$$

As an example, suppose the Hyatt Hotel described above estimates that the number of customers who book a room but fail to show up on the night in question is Normally distributed with mean 20 and standard deviation 10. Moreover, Hyatt estimates that it costs $300 to "bump" a customer (Hyatt receives no revenue from this customer, and the $300 is the cost of alternative accommodation plus a possible gift certificate for a one-night stay at Hyatt in the future). On the other hand, if a room is not sold then Hyatt loses revenue equal to one night sold at a discount since at the very least this room could have been sold to leisure customers[6]. How many bookings should Hyatt allow? The critical ratio is

$$\frac{B}{B+C} = \frac{105}{300+105} = 0.2592.$$

From the Normal distribution table we get $\Phi(-0.65) = 0.2578$ and $\Phi(-0.64) = 0.2611$, so the optimal optimal *z** is approximately *-0.645* and the optimal number of rooms to overbook is *Y*=20-0.645*10=13.5* (one can also use Excel's *Norminv* function for this calculation: "*=Norminv(0.2592,20,10)*" gives us an answer of *13.5*). If we round up to 14, this means that up to *210+14=224* bookings should be allowed.

7. Complications and Extensions

There are a wide variety of complications we face when implementing a yield management system. Here we discuss a few of the more significant challenges.

<u>Demand Forecasting</u>

In the examples above, we used historical demand to predict future demand. In an actual application, we may use more elaborate models to generate demand forecasts that take into account a variety of predictable events, such as the day of the week, seasonality, and special events such as holidays. In some industries, such as retail clothing, greater weight is given to the most recent demand patterns since customer preferences change rapidly. Another natural problem that arises during demand forecasting is censored data, i.e., company often does not record demand from customers who were denied a reservation.

In our example in Sections 3-5 used a discrete, empirical distribution to determine the protection level. A statistical forecasting model would generate a continuous distribution, such as a Normal or *t* distribution. Given a theoretical distribution and its parameters, such as the mean and variance, we would again place the protection level where the distribution has a cumulative probability equal to the critical fractile.

Variation and Mobility of Capacity

Up to now, we have assumed that all units of capacity are the same; in our Hyatt example we assumed that all 210 hotel rooms were identical. However, rooms often vary in size and amenities. Airlines usually offer coach and first-classes. Car rental firms offer subcompact, compact, and luxury cars. In addition, car rental firms have the opportunity to move capacity among locations to accommodate surges in demand, particularly when a central office manages the regional allocation of cars. The EMSR framework described above can sometimes be adapted for these cases, but the calculations are much more complex. Solving such problems is an area of active research in the operations community.

Customers in a Fare Class Are Nor All Alike

While two leisure travelers may be willing to pay the same price for a particular night's stay, one may be staying for just one day while the other may occupy the room for a week. A business traveler on an airplane flight may book a ticket on just one leg or may be continuing on multiple legs. Not selling a ticket to the latter passenger means that revenue from all flight legs will be lost.

In each of these cases, the *total* revenue generated by the customer should be incorporated into the yield management calculation, not just the revenue generated by a single night's stay or a single flight leg. There are additional complications when code-sharing partners (distinct airlines that offer connecting flights among one another) operate these flight legs.

---

[6] Technically, this empty room could have been sold at either low or high price so it is possible that lost revenue is actually higher than we suggest. Also, often hotels charge certain penalty for late cancellations.

If code-sharing occurs, then each of the partners must have an incentive to take into consideration the other partners' revenue streams.

A similar complication occurs when there are group bookings. How should the hotel consider a group booking request for 210 rooms at a discount rate? Clearly, this is different from booking 210 individual rooms since denying a reservation to one out of 210 group customers may mean that all 210 will be lost to the hotel. Frequently, companies restrict group reservations during the peak seasons but up till now there are no general rules for handling group reservations.

<u>Summary</u>

Table 2 summarizes the impact of these issues on three industries: airlines, hotels and car rental firms[7]. Sophisticated yield management tools have been developed in all three industries, and these tools take industry-specific factors into account. However, all of these tools are based on the basic EMSR model described above.

Parameter	**Airline**	**Hotel**	**Car rental**
Unit of capacity	Seat	Room	Car
Number of resource types	2-3 (e.g., 1st-class and coach seats)	2-10+	5-20+
"Capacity" at a location fixed or variable	Fixed	Fixed	Variable
Mobility of capacity	Small	None	Considerable
Number of possible prices per unit	Many (3-7+)	Few (2-3+)	Many(4-20+)
Duration of use	Fixed	Variable	Variable
Corporate discounts	Occasional	Yes	Yes
Capacity managed locally or centrally	Central	Central/local	Central/regional/local

Table 2. Comparison of yield management applications.

[7] Adopted from Carrol W. J. and R. C. Grimes. Evolutionary change in product management: experiences in the car rental industry. Interfaces, 1995, vol.25, no.5, 84-104.