PART 4

Matching Supply and Demand

The management of capacity in service firms frequently presents a more difficult and more expensive problem than in manufacturing. To determine capacity in many manufacturing firms, one usually considers long-run average demand. Inventory is frequently used to compensate for annual seasonality or short-term fluctuations in demand. In services, however, although the long-run average must be considered, the short term is vital as well. Because of the perishable nature of service sector “inventory” such as airplane seats or hotel rooms, capacity plans must consider demand by day of the week and even time of day—a level of detail that manufacturers usually find irrelevant.

Further, capacity planning mistakes are often more costly in services. When caught short on capacity many manufacturers simply “back-order” a customer’s request. In many services, though, back orders simply cannot exist—quite literally, “the ship has sailed” for a cruise operator. Not having what the customer wants when she arrives can mean either a one-time lost sale to the competition, or possibly the defection of the entire stream of future sales of that customer and whoever else she chooses to tell about her experience.

Many strategic and tactical decisions must be made concerning services capacity. A “one-size-fits-all” strategy will not work. Even though it may be in the strategic interest of one airline to tightly match capacity to demand by heavily overbooking flights, another airline might be most profitable flying at far below capacity on average. Many of the qualitative aspects of these topics were covered in Chapter 2. This part of the book looks at putting those general strategies into action.

CHAPTER

12

Yield Management

LEARNING OBJECTIVES *The material in this chapter prepares students to:*

- Understand the need for overbooking.
- Use three different methods to calculate an overbooking level.
- Determine how to allocate service capacity among customer groups.
- Understand the intricacies of pricing for a capacity constrained service.

Capacity Strategies

Capacity planning for many service firms can be far more difficult than for manufacturers. Manufacturers can set capacity by looking at long-run average demand. For many service firms, however, long-run averages become somewhat meaningless when capacity must react to general seasonality, daily demand variations, and time-of-day demand fluctuations. If the average manufacturer found out that most end consumers bought its product between 2 P.M. and 3 P.M., this knowledge wouldn't change its capacity strategy at all, but it would be important information for many service firms.

Capacity decisions in service firms are not only more complex than in manufacturers, but can be more important as well. Manufacturers deal with short-term imbalances in production and demand by either carrying inventory or creating a backorder list for later shipment. In most services, the "inventory" of capacity is employee time, or a fixed asset not being used, such as a hotel room or an airplane seat, so excess inventory cannot be stored for later use. Backorders quite often cannot occur: Imagine a sales clerk at a department store stating that he will be able to speak with a customer by next Tuesday. Consequently, a temporary imbalance in supply and demand can result in either idle employees and resources if demand is smaller than supply, or lost sales to the competition if demand is larger than supply. (Service firms that can use physical inventory are discussed in Chapter 13.)

These factors turn simple tactical decisions into strategic ones. Consider this simple example of the basic strategic direction for service capacity. An ice cream parlor experiences the following demand for ice cream cones:

Weekdays	100 – 300
Saturday	500 – 1,500
Sunday	500 – 1,100

For the manufacturer supplying the cones, capacity is a simple matter: It calculates average weekly demand:

$$5(200) + 1{,}000 + 800 = 2{,}800 \text{ cones}$$

It makes 2,800/7 = 400 cones every day, and carries a small inventory of extra cones for the busier days. For the service provider who fills cones as customers walk in, however, simple arithmetic no longer applies. A strategic decision must be made. The ice cream parlor manager may use one of the four basic strategies outlined next.[1] When considering these strategies, assume that one employee can make 100 cones per day.

1. **Provide: Ensure sufficient capacity at all times.** To carry out a provide strategy, one would want to always have enough people to handle the maximum demand, so the firm would have 15 employees working on Saturday, 11 on Sunday, and 3 the rest of the week. It is usually difficult to employ significant numbers of part-timers, so this strategy would employ enough full-time employees to meet those numbers. This strategy is associated with a high-service quality generic strategy, but it is also high cost, and would result in significant idle time for employees. Businesses with these characteristics include high-margin sales (e.g., jewelry, luxury automobiles) and those with wealthy individuals as clients (e.g., chauffeuring, private banking). Also, firms that compete on delivery speed (often called "time-based" competitors) should adopt this approach.
2. **Match: Change capacity as needed.** This strategy would use ten employees on Saturday, eight on Sunday, and two the rest of the week, with the excess Saturday and Sunday employees strictly part-timers. This approach balances service quality and costs and is representative of a large number of firms, including most mid- and low-priced restaurants and telemarketing firms.
3. **Influence: Alter demand patterns to fit firm capacity.** Here, pricing, marketing, or appointment systems flatten demand peaks to conform to capacity. It is most common in high capital-intensive services such as airlines and hotels, but highly paid professionals such as medical doctors and lawyers also commonly use it.
4. **Control: Maximize capacity utilization.** If only full-time employees could be used, five days per week, this strategy would have just two employees whose schedules overlapped on weekends. The generic strategy behind this option is to compete on cost by driving employee idle time to zero. It is often used in the public sector (e.g., driver's license bureaus) and low-margin services, as well as situations where high-priced employees want to maximize their utilization. Many physicians deliberately schedule patient appointments so tightly that a crowd is always in their waiting room. This strategy is willing to sacrifice sales at busy times to ensure the service functions efficiently all the time.

To assist in crafting these strategies, a host of specific tactics can be used to manage supply and demand (an in-depth discussion of these issues can be found in Klassen and Rohleder, 2001). Supply management tactics include the following:

1. Crandall and Markland (1996).

- **Workshift scheduling.** The unevenness of customer demand throughout a day means utilizing creative work schedules, such as nonuniform starting times, and workdays that have variable work hours. Work scheduling software is available to help construct flexible solutions within a match strategy.
- **Increasing customer participation.** A traditional method for a control strategy cuts total labor by encouraging customers to participate in serving themselves. For example, many fast-food restaurants use a semi-control strategy in which customers pour their own fountain drinks and procure their own condiments.
- **Adjustable (surge) capacity.** "Surge" capacity means capacity that can be available for a short period of time. By cross-training personnel for different jobs, a company can flexibly shift personnel temporarily to increase the capacity of any one position. Because cross-training is expensive to undertake, and cross-trained personnel are more expensive to retain, it is an appropriate approach within a provide strategy.
- **Sharing capacity.** Capacity can often be shared between departments or between firms for personnel or equipment that is needed only occasionally. For example, small business incubators often contract with dozens of businesses to share the same secretarial, accounting, and office management team.

Several tactics can be used to manage demand as well.

- **Partitioning demand.** It is not unusual for some components of demand to be inherently random, while some are fixed. This approach melds the more malleable demand around the tendencies of the random demand. That is, if it is known that more walk-in business generally comes in from 11 A.M. to 1 P.M., then schedule appointments either before or after that time. This approach works primarily for provide and match strategies.
- **Price incentives and promotion of off-peak demand.** This highly common method works in an influence strategy, which many of us see in our telephone bills. It is also commonly used in restaurants ("early bird" specials), hotels (both off-season and day-of-week pricing), resorts, and so on.
- **Develop complementary services.** The way to avoid the inevitable seasonality of many services is to couple countercyclical services together: Heating and air conditioning repair, ski slopes in winter and mountain bike trails in the summer. Unfortunately, this approach remains only a theoretical construct for most services.
- **Yield management.** Yield management combines three techniques: (1) overbooking, (2) assigning capacity amounts to different market segments, and (3) differential pricing in different market segments. It is used extensively by many industries and is the subject of the remainder of this chapter.

Yield Management

Consumers encounter examples of what is called *yield management* constantly. A little knowledge about how these systems work can make life easier, or at least less expensive. Some practical examples of dealing with a yield management system include overbooking at a car rental agency. Even though you "confirmed" your reservation, it still pays to show up early in the day to get a car; sometimes those who show up late are out of luck. If you are a little more flexible about which days you fly, an airplane ticket may cost several hundred dollars less. The airline flight you are trying to book a seat on may be full today, but be patient and keep trying; tomorrow a seat may be available, even without others' cancellations. A hotel that says no room is available for you on Thursday night may suddenly find a room for Thursday if you add that you are staying Friday.

These situations occur because of yield management systems. The application of yield management practices often leave customers and employees puzzled. This chapter introduces the reasoning and techniques of yield management. Even if you do not work in an industry where yield management is practiced, this material will at least help you be a better consumer.

The term *yield management* itself is a bit of a misnomer because these techniques are not directly concerned with managing "yield" but are really concerned with managing revenue. Consequently, the set of techniques described in this chapter is sometimes called *revenue management* or *perishable asset revenue management*.

The purpose of yield management techniques is to sell the right capacity to the right customer at the right price. Not every firm can use these techniques, but many capital-intensive services can and do use them heavily. The main business requirement for using the techniques of this chapter is having limited, fixed capacity. Many other business characteristics make yield management more effective:

- Ability to segment markets
- Perishable inventory
- Advance sales
- Fluctuating demand
- Accurate, detailed information systems

These characteristics increase the complexity of a business and the profit potential from applying yield management.

Industries that currently fully utilize yield management techniques are transportation-oriented industries, such as airlines, railroads, car rental agencies, and shipping; vacation-oriented industries, such as tour operators, cruise ships, and resorts; and other capacity-constrained industries, such as hotels, medicine, storage facilities, and broadcasting (selling commercial time). Many other industries can partially use these techniques.

Yield management is a relatively young science. Airlines are credited with the invention of most of these techniques, especially Sabre, formerly with American Airlines (see the Service Operations Management Practices: Yield Management Increases Revenue $1 Billion/Year at American Airlines). However, the airlines did not develop most of these systems until a few years after the industry was deregulated in 1978, and the techniques only began to spread to other industries in the 1990s.

A yield management system consists of three basic elements:

1. Overbooking (accepting more requests for service than can be provided)
2. Differential pricing to different customer groups
3. Capacity allocation among customer groups

Each of these elements will be discussed in turn, then some practical implementation issues will be addressed.

OVERBOOKING

The need for overbooking is clear. Customers are fickle and do not always show up, so firms that overbook make far more money than those that don't. American Airlines

SERVICE OPERATIONS MANAGEMENT PRACTICES

Yield Management Increases Revenue $1 Billion/Year at American Airlines

It can be challenging to turn a profit in the airline business, with margins usually in the 1% to 5% range. Once flights are scheduled, costs are essentially fixed, and they can only hope to fill part of the plane with customers who aren't as fussy about price. Yield management got its start at American Airlines when the industry was deregulated. New startups PeopleExpress and World Airways were offering one-way fares from New York to San Francisco for $99—less than half the regulated fare. With their higher operating costs, the traditional airlines like American couldn't exist at such prices, because even a full load of $99 passengers would mean losses. Through its Sabre unit, American responded by inventing yield management. They matched the low fares, but allowed only a portion of their planes to be filled by them, while the newcomers sold every seat for the same price.

In a few years most of the upstarts were out of business, and yield management became more sophisticated. The founder of PeopleExpress, Don Burr, claimed that the superior yield management abilities of their competitors caused their demise (Cross, 1997, p.125). American estimates that its yield management system currently adds $1 billion per year in revenue. For American, whose annual profits are rarely above that figure, yield management is the difference between profitability and bankruptcy.

The CEO of American Airlines, Bob Crandall, said that "Yield management is the single most important technical development in transportation management since . . . deregulation" (Cross, 1997, p.127).

Source: Adapted from Cook (1998).

estimated that their overbooking system garners them an additional $225 million in profit annually (Smith, Leimkuhler, and Darrow, 1992).

If airlines did not overbook, planes that are now full would fly an average of 15% empty. "No-shows cost the world's airlines $3 billion annually, even after efforts to minimize the revenue loss by overbooking" (Cross, 1997, p. 146). No-shows for restaurant reservations average about 10%, with some reporting 40% no-shows during the Christmas holidays. It has been reported that rental car no-shows in the Florida market reached 70% of reservations. Of course, the alternative to overbooking is to simply charge the customers whether they show up or not. Unfortunately, that approach failed in restaurants and auto rental businesses, and other businesses discarded it out of hand. Consumer resistance was high: Imagine missing your plane flight due to traffic, only to be told, "The seat you paid for is in the sky, the ticket you have is worthless. The next flight out will cost you another $500, even though they have empty seats on that one." So the question for many businesses is not *whether* to overbook, but rather, *how much* to overbook.

To demonstrate some mathematical methods to help determine the level of overbooking, consider the following example.

EXAMPLE 12.1: *The Hotel California*

The Hotel California found that it frequently turned down a customer in the lobby because a room was reserved for a customer who never showed up. The manager, felt that the hotel's policy of overbooking should be examined.

The average room rate was $50 per night, but the hotel could not collect the room rate from the no-show customers. If no overbookings were allowed, each no-show would in reality cost the hotel $50. If it overbooked too much and filled up early in the night, customers with reservations who arrived later to find no rooms available would be most unhappy. About 10% of those customers did not cost the hotel any money; they merely muttered menacingly and walked out. Another 10% were satisfied with being "walked" (or transferred) to another hotel at no cost to the Hotel California. The remaining guests were so upset by this situation that the hotel had to repair broken lobby furniture at a cost of $150.

The hotel's no-show experience is summarized in Table 12.1. What should the overbooking policy be? We will discuss three approaches to answering that question.

Overbooking Approach 1: Using Averages

In Table 12.1 the average number of no-shows is calculated by

$$0(0.05) + 1(0.10) + 2(0.20) + 3(0.15) + \ldots + 10(0.05) = 4.05.$$

Since the average number of no-shows is four, it might seem reasonable to take up to four overbookings.

This approach offers the advantages of being intuitive and easy to explain. It is also usually better than doing no overbooking at all. It fails, however, to weigh the relevant costs, which presents a significant disadvantage. For instance, if the cost of a disgruntled customer is nothing, then the best policy would be to overbook 10 every night to ensure that the hotel is full. That is, if all the customers who had reservations and didn't get rooms simply left at no cost to the hotel, the hotel would just be concerned about losing the potential $50 of a paying guest. Likewise, if all the disappointed customers reacted by telling Norman's mother on him—the equivalent of an infinite cost to Mr. Bates—Norman would never overbook.

Overbooking Approach 2: Spreadsheet Analysis

The two costs to consider here are:

C_o = Overage (customers denied advance reservation with rooms left unoccupied, often called "spoilage" in industry)

C_s = Stockouts (customers with reservations are turned away because no rooms are left, called "walked" customers in the hotel industry and "spill" by the airlines)

In this case, C_o = $50, the cost of the room, and C_s = 0.2($0) + 0.8($150) = $120.

TABLE 12.1: *Hotel California No-Show Experience*

No-Shows	% of Experiences	Cumulative % of Experience
0	5	5
1	10	15
2	20	35
3	15	50
4	15	65
5	10	75
6	5	80
7	5	85
8	5	90
9	5	95
19	5	100

One way to put the relevant costs into the picture is to use the spreadsheet shown on Table 12.2 (This spreadsheet is also on the CD included with this text.)

This spreadsheet calculates the expected cost for every possible scenario. For example, if no overbooking is done, then the column labeled "0" shows that on the 5% of days when there are zero no-shows, there's no cost at all, but on the 10% of days when there is one no-show, the cost is $50. The total cost at the bottom sums up 0.05($0) + 0.10($50) + . . . + 0.05($500) = $203. The overbooking level with the lowest expected cost is to overbook two rooms, with an expected cost of $137.

The advantages of this method are that it incorporates relevant costs and can be spreadsheet based and fairly easy to figure out. Also, as will be seen shortly, if the costs and revenues are uncertain or not quite as easy to figure out as in the Hotel California example, then this method can be readily adapted. Two disadvantages of this method, though, are that it requires accurate data and it is a "brute force" type of technique that does not increase a manager's intuition about the problem.

Access your Student CD now for Table 12.2 as an Excel spreadsheet.

TABLE 12.2: *Hotel California Overbooking Cost*

		Number of Reservations Overbooked										
No-Shows	**Probability**	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
0	0.05	$ 0	$120	$240	$360	$480	$600	$720	$840	$960	$1,080	$1,200
1	0.10	$ 50	$ 0	$120	$240	$360	$480	$600	$720	$840	$ 960	$1,080
2	0.20	$100	$ 50	$ 0	$120	$240	$360	$480	$600	$720	$ 840	$ 960
3	0.15	$150	$100	$ 50	$ 0	$120	$240	$360	$480	$600	$ 720	$ 840
4	0.15	$200	$150	$100	$ 50	$ 0	$120	$240	$360	$480	$ 600	$ 720
5	0.10	$250	$200	$150	$100	$ 50	$ 0	$120	$240	$360	$ 480	$ 600
6	0.05	$300	$250	$200	$150	$100	$ 50	$ 0	$120	$240	$ 360	$ 480
7	0.05	$350	$300	$250	$200	$150	$100	$ 50	$ 0	$120	$ 240	$ 360
8	0.05	$400	$350	$300	$250	$200	$150	$100	$ 50	$ 0	$ 120	$ 240
9	0.05	$450	$400	$350	$300	$250	$200	$150	$100	$ 50	$ 0	$ 120
10	0.05	$500	$450	$400	$350	$300	$250	$200	$150	$100	$ 50	$ 0
Total Cost		$203	$161	$137	$146	$181	$242	$319	$405	$500	$ 603	$ 714

Overbooking Approach 3: Marginal Cost Approach

Using a little algebra, this method comes at the problem mathematically by noting that one would like to keep accepting bookings until the expected revenue is less than or equal to the expected loss from the last booking. Mathematically, increase bookings until

$$E(\text{revenue of next booking}) \leq E(\text{cost of next booking})$$

which is the same as

Revenue of filling a room × Probability of more no-shows than overbooked rooms ≤ Cost of dissatisfied customer × Probability of fewer or the same number of no-shows than overbooked rooms

Or, in the mathematical terms used previously,

$$C_o \times P(\text{Overbookings} < \text{No-shows}) \leq C_s \times P(\text{Overbookings} \geq \text{No-shows})$$

which can be converted to

$$C_o \times [1 - P(\text{Overbookings} \geq \text{No-shows})] \leq C_s \times P(\text{Overbookings} \geq \text{No-shows})$$

or equivalently,

$$C_o - C_o \times P(\text{Overbookings} \geq \text{No-shows}) \leq C_s \times P(\text{Overbookings} \geq \text{No-shows})$$

Adding $C_o \times P(\text{Overbookings} \geq \text{No-shows})$ to both sides and dividing both sides by $(C_o + C_s)$ leaves the basic overbooking formula: Accept bookings until

$$C_o / (C_s + C_o) \leq P(\text{Overbookings} \geq \text{No-shows}) \qquad (12.1)$$

In the preceding problem, this calculation leads to

$$\$50/(\$120 + \$50) = 0.29$$

Looking at Table 12.1, the smallest number of overbookings at which $P(\text{Overbookings} \geq \text{No-shows})$ is 2, where the cumulative probability of no-shows reaching this level is 0.35.

This basic formula is easy to remember and apply, even to informal data. For example, C_o, the lost potential revenue, may be easy to figure in most circumstances, but C_s is not, and usually must be estimated. Also, the cumulative probability distribution of no-shows is often not accurately known. So a general feel that, say, a complaining customer is three times as costly as the potential revenue means that a manager would only want to overbook until $P(\text{Overbookings} \geq \text{No-shows})$ is about $1/(1 + 3) = 25\%$. So, if the average number of no-shows is about 15 with a standard deviation of five, using the traditional z-score calculations from standard statistics texts, about 12 overbookings might be appropriate.

Although this formula is simple to use, it presents a significant drawback. Equation (12.1) implicitly assumes a linear cost of dissatisfied reservation holders; that is, if only one customer in your hotel lobby or airport lounge is dissatisfied and will cost \$300 to placate, then 20 dissatisfied customers will cost 20 × \$300 = \$6,000 to satisfy. Unfortunately, that answer is not always the case. As shown in Figure 12.1, the cost curve for overbooking can increase per person with the number of unhappy customers. A roomful of 20 unhappy customers is far more of a problem than 20 instances of a single unhappy customer. Although this formula does not account for this contingency, the spreadsheet method can easily be programmed with it in mind.

FIGURE 12.1: *Actual Versus Linear Overbooking Cost Curve*

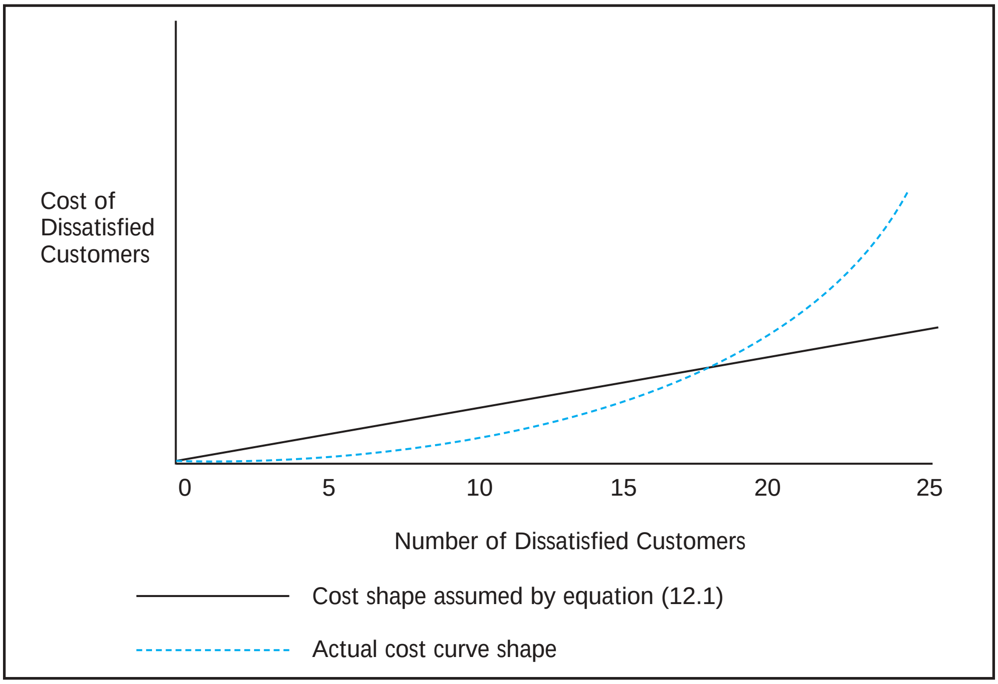

Dynamic Overbooking

The overbooking decision is often not a one-time, "static" decision. Rather, it is a decision that is "dynamic," which here just means that it changes over time. In a typical situation for a firm that takes reservations a long time in advance, the dynamic overbooking curve is shaped like Figure 12.2.

When the event is still a long time in the future, the allowed number of overbookings for it is at a peak. As the event nears, the number of allowed overbookings drops. This practice reaches its logical conclusion at the time the event takes place, when allowed overbookings drop to zero. For example, if three people are running to the plane for the final boarding and only two seats are available, the third person will not be sold a seat that cannot possibly be used.

Allocating Capacity

A difficult problem that afflicts many firms is allocating capacity among their customer groups. That is, when to say "no, we're full" to one customer while holding open capacity in hopes that a more profitable customer will arrive later. For airlines and hotels, especially, reservation activity follows Figure 12.3: High-revenue customers tend to make reservations fairly close to the event, while low-revenue customers often make reservations months in advance. In the airline industry, the price-conscious vacationers make reservations months in advance, while the high-paying business travelers may make reservations close to the event. For hotels, price-conscious group business may make reservations a year in advance, while more highly profitable transient business may simply walk in the door that day. For the sake of simplicity, this chapter will focus on segmenting capacity between just two

FIGURE 12.2: *Dynamic Overbooking*

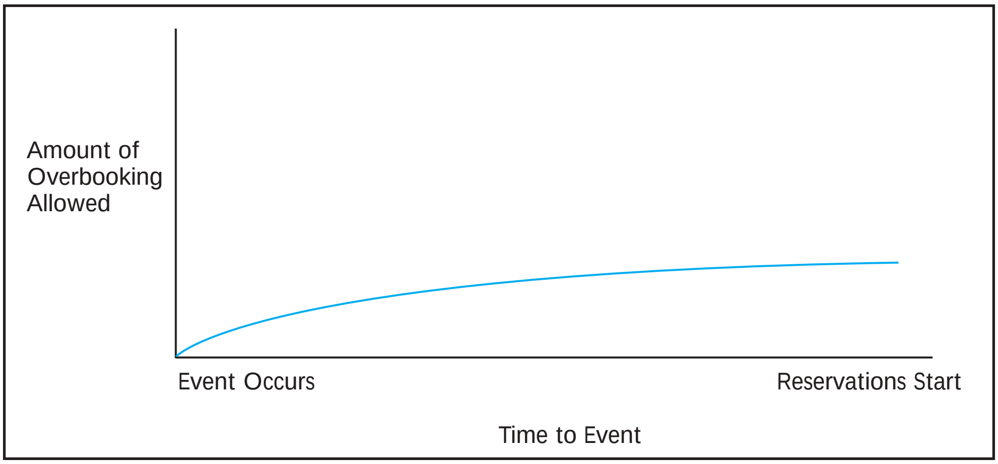

customer groups. The real problem, however, is far more complex. A modern airline can have 10 different customer segments on the plane, each of which requires a capacity allocation decision (Cross, 1004).

The situation depicted in Figure 12.3 is complex because one cannot simply let all reservations be taken first-come, first-served. Doing so would cause a firm to turn away a substantial portion of their most profitable customers while filling up on the less profitable ones. Also, most firms cannot simply say "no" to lower-revenue business, because they cannot fill their capacity solely with high-revenue business. Consequently, the firm must decide ahead of time at what point to shut off the low-revenue business in anticipation of the high-revenue business booking later.

FIGURE 12.3 *Cumulative Reservation Activity*

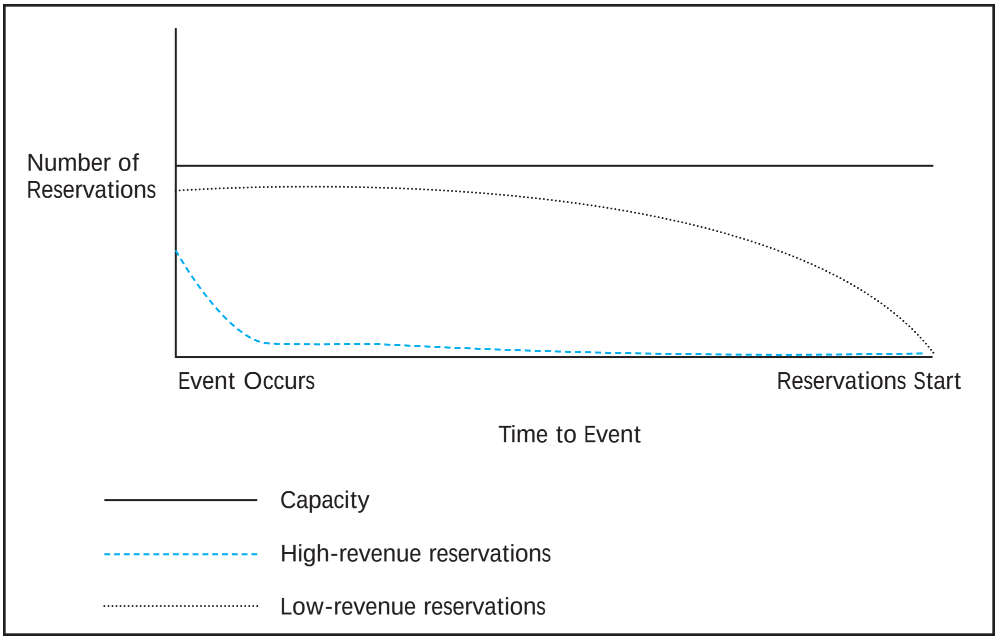

The general methods used to solve this problem can be classified as follows:

- Nested versus distinct
- Static versus dynamic

To demonstrate the differences in these methods, let's first describe an example problem.

EXAMPLE 12.2: *Chancey Travel*

Chancey Travel is offering a cruise to the Antarctic and wants to fill the 100 cabins available. Its primary market is Premium customers. Premium customers pay $12,000 for the trip, and the variable costs of serving these customers amounts to $2,000 per trip, as they are pampered endlessly and plied with Godiva chocolate three times a day. According to market analysis, the demand from Premium passengers is uniformly distributed between 51 and 100, which means a 2% chance that 51 Premiums will sign up, a 2% chance that 52 will go, and so on up to a 2% chance that the entire boat can be filled with 100 Premiums. This means that if the entire boat were reserved for Premiums an average of 75 Premiums would be on board. Premiums come from the idle rich, who tend to make their decisions at the last minute.

Because the probability that the entire boat cannot be filled with Premiums is substantial, another market is sought. Market research discovered customer interest in a "rough" experience to Antarctica, but these consumers were more price sensitive. These Discount customers paid $2,500 for the trip, but the marginal cost of serving them was $0: heat was turned off to their cabins (hot water bottles were available for a proper deposit), and they received no food (they had to catch penguins with their own equipment). Demand was such that the ship could be entirely filled with Discounts, and these customers were willing to book far in advance.

The dilemma for Chancey Travel is: How many cabins should be reserved in hopes of Premium customers? Several methods for arriving at an answer are discussed here.

Capacity Allocation Approach 1: Static Methods: Fixed Number, Fixed Time Rules

We begin with static or one-time decision rules. Basic fixed number and fixed time rules are the easiest to implement. A fixed time rule means simply that a firm will accept discount bookings until a specific date. No limit is set on the number of discount sales. The motivation for this type of rule is the transparency to the customer and the ease of implementation, but it is clearly not close to being optimal. A fixed number rule for the problem may be to, say, reserve exactly 75 slots for Premium customers and exactly 25 slots for Discount customers. The amounts reserved for each group can be viewed as a quota. Although a step up from the fixed time rule, it too presents certain problems. For example, if 75 Premiums and 20 Discounts are currently signed up, and a Premium customer wants to pay, this type of rule states that the Premium customer must be turned away because the Premium slots are already filled.

Capacity Allocation Approach 2: Nested Static Methods

To get around the somewhat absurd problem of turning down more profitable customers in hopes that less profitable ones will appear later, so-called nested methods are used. Let's still assume that, because an average of 75 Premium bookings is expected, we want to reserve 75 slots for Premiums. Instead of the remaining 25 slots being given to Discounts, the remaining 25 are sold on a first-come, first-served basis. In this manner, the 75 slots for Premiums can be thought of as a "protection level;" that is, at least 75 rooms are protected just for Premiums, but they can get more. In this nested system other groups are not allowed into the "nest" of 75 protected rooms, but the protected group can venture outside that number.

Our use of 75 as a hypothetical number for the rooms protected for the Premiums does not offer the best strategy, however. To show how some basic calculations can lead to a better strategy, the expected marginal seat revenue (EMSR) heuristic is introduced here (Belobaba, 1989). The EMSR heuristic provides a basic logic for determining how much protection to give different customer levels. The EMSR heuristic is no longer widely used in industry. However, the methods that are used currently in industry are both highly complex and proprietary.

The EMSR heuristic allocates capacity one unit at a time. For this problem, allocating the 1st through 51st rooms is fairly simple—they should all go to Premiums, since we are certain that at least 51 Premiums will show up. For deciding whether to protect the 52nd room we compare the expected marginal revenue from each group and assign the room to the group that provides the most expected revenue. For allocating the 52nd room, we are 98% certain a Premium customer will request it, so the calculations are as follows:

$$\text{Premium: } 98\%(\$12{,}000 - \$2{,}000) = \$9{,}800 \text{ expected revenue}$$

$$\text{Discount: } 100\%(\$2{,}500) = \$2{,}500 \text{ expected revenue}$$

Because $9,800 > $2,500, the room is reserved for Premiums. For allocating the 53rd room, the following calculation applies:

$$\text{Premium: } 96\%(\$10{,}000) = \$9{,}600 \text{ expected revenue}$$

$$\text{Discount: } \$2{,}500 \text{ expected revenue}$$

Again the room is reserved for Premiums, because $9,600 > $2,500. This process continues until we reach the 88th seat:

$$\text{Premium: } 24\%\ (\$10{,}000) = \$2{,}400 \text{ expected revenue}$$

$$\text{Discount: } \$2{,}500 \text{ expected revenue}$$

Because $2,400 < $2,500, the protection level for Premiums stops at 87. The EMSR heuristic therefore states that we should allocate 87 rooms Premium, 13 rooms Discount. Note, however, that on an average boat, this allocation would result in 75 Premium passengers, 13 Discount passengers, and 12 empty rooms. This result explains why, when one is allocating capacity as best as one can, it frequently results in unused capacity.

This procedure can be summarized in a rule similar to the overbooking formula given in equation (12.1) earlier. Set the protection level of Premiums at the smallest number where:

$$P(\text{Premium demand} \geq \text{Protection level}) \leq \text{Discount revenue/Premium revenue} \quad (12.2)$$

Here, Discount revenue/Premium revenue = $2,500/($12,000 – $2,000) = 0.25, and there is a 26% chance that Premium demand is greater than 87.

To see how to use this from a different perspective, consider Premium demand for Chancey Travel to be normally distributed, with a mean of 70 and a standard deviation of 20. Let's make Discount revenue $3,000, instead of $2,500. So, Discount revenue/Premium revenue = 0.3, and we are searching for the 30th percentile of demand.

This is a one-tailed look-up of an "area of the standard normal distribution" table (located at the back of the book). So one desires to look up 0.5 − 0.3 or 0.2 in the body of the table, which yields a z-score of 0.515. Applying this z-score to the problem data, mean = 70, std dev = 20, $70 + 0.515 \times 20 = 80$ seats should be protected for the Premium passengers.

This approach is not limited to just two customer categories. The protection level for each customer category can be calculated by comparing the revenue associated with each successive customer class.

Capacity Allocation Approach 3: Dynamic Methods

The previous methods describe one-time decisions made in advance. However, in many cases, a better but more complicated method is available. The probability distribution used previously (uniform distribution between 51 and 100) is a forecast. But forecasts become more accurate as the event gets closer in time, and early activity tends to be a good predictor of what will eventually unfold.

For example, consider Table 12.3, which tracks the applications to a well-known Southern MBA program. Although only a small percentage of total applications arrive before December 1, the amount received by that date is a good predictor of the eventual number.

Applying that notion to yield management means that the type of analysis done with the EMSR heuristic earlier is constantly reassessed. In Figure 12.4, if we are continually forecasting demand from early reservations history, reservation activity 35 days before the event may look as though we are on the "low" demand curve for high-revenue customers. Consequently, the floodgates are opened up to lower-revenue customers. However, 30 days before the event, reservation activity looks as though a "medium" number of high-revenue customers will be forthcoming, so the protection level for high-revenue customers is increased. Finally, Figure 12.4 shows that 20 days before the event, high-revenue customer class reservation activity is such that the entire service capacity can be taken up with high-revenue customers, so all lower-revenue customer classes are cut off from further reservation activity. Of course, this scenario can happen in reverse order as well, with more and more capacity being allocated over time to lower-revenue customer classes. Consequently, as a discount-seeking consumer, one may find a flight to be "full" one day, yet be able to get a reservation the next day for the same price due to this constant reevaluation.

To be done accurately, this type of nonstop capacity allocation requires immense amounts of data regarding customer behavior. For example, Delta Air Lines forecasts passenger demand on 16.5 million future combined flight/customer classifications

TABLE 12.3: *MBA Program Applications*

	Number of Applications Received By:	
Year	**December 1**	**June 1**
1	57	852
2	89	931
3	110	1,023

FIGURE 12.4: *Dynamic Capacity Allocation*

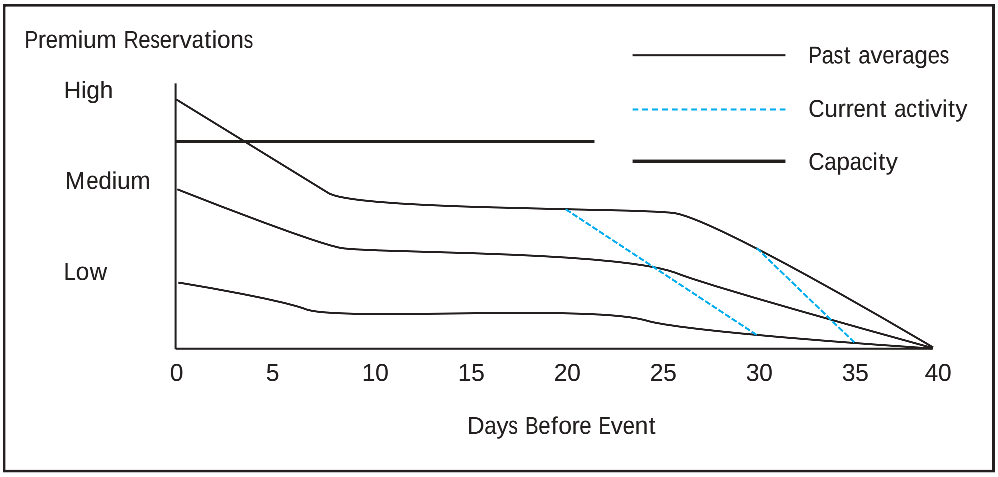

every day. The database for this forecast requires 2 terabytes of data (Cross, 2004). Unfortunately, such data do not exist for many firms. Also, events unknown to any computer model may cause reservation activity to suddenly drop or pick up, such as competitor pricing changes, new competitors coming on line, or such things as sporting events and conventions being scheduled where they did not exist before. Because of these difficulties, computer models help, but the final decisions on capacity allocation for many industries is still manual, performed by managers with an intuitive sense of the risks.

Complexities of Capacity Allocation

So far we considered allocating capacity for single events. However, the situation is often more complex. For example, consider a low-revenue customer booking a hotel room for a Thursday night versus booking for both Thursday and Friday nights. If the hotel has mainly a business clientele, then Thursdays tend to be fairly full of high-revenue customers, but on Friday night the hotel is relatively empty. Consequently, allowing a low-revenue customer to take up a valuable Thursday room might be unprofitable, but if the low-revenue customer also uses a Friday room that would ordinarily go unoccupied, the profitability changes. Consequently, the protection level of the high-revenue customer class for a Thursday night depends on the lower-revenue customer length of stay.

In the jargon of the hotel yield management industry, a "bid price" strategy is used to decide whether to accept such reservations; that is, the actual price of the lower-revenue customer is compared to the expected revenue over the course of the proposed length of stay.

Pricing

Thus far we assumed prices are exogenous, or outside of our control, which of course is not the case. This section discusses the difficult topic of setting prices in a yield management system.

The general idea in yield management is to break up a market into a number of different customer segments and charge different prices for each segment. Figure 12.5 shows the traditional supply and demand curves seen in every introductory economics class.

FIGURE 12.5: *Traditional Supply and Demand Equilibrium*

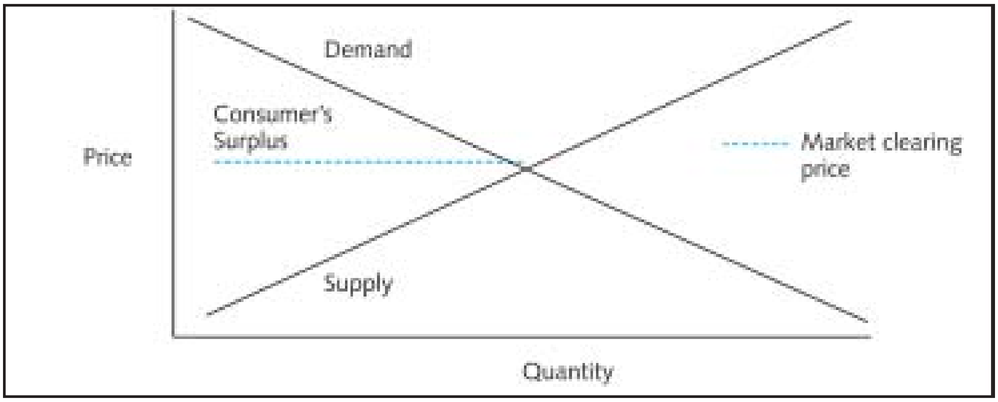

The market-clearing price creates "consumer's surplus" for customers who would have been willing to pay more and economic rent for suppliers whose costs are lower. The idea of yield management is to segment different customer categories with separate market clearing prices, so that high prices are charged to those who would be willing to pay them and low prices are charged to those who would ordinarily not use the service if the single market clearing price were in place. Figure 12.6 shows a general representation of three markets for airline seats.

Customer class 1 (first-class passengers) wants premium service and flexibility, and cares little about cost; customer class 2 (business class) values the flexibility of making last-minute changes over price; customer class 3 (often vacationing families) makes plans well in advance and are highly sensitive to price. The consumer surplus from each group is the area formed by the triangle of the market price (dotted line), demand curve, and the vertical line extending from the dotted line to the customer class bar. Two important economic effects come into play here: Large portions of consumer surplus are shifted to the supplier, and bringing on customers who would not ordinarily use the service can expand the market itself.

The magnitude of pricing differences can be enormous. Table 12.4 shows a snapshot of some airline fares. Here, premium first-class and business class are not considered—just full coach fare versus discounts on that same coach seat—and the cheapest prices are 11% to 13% of the cost of a full coach fare.

FIGURE 12.6: *Supply and Demand Equilibrium in Yield Management*

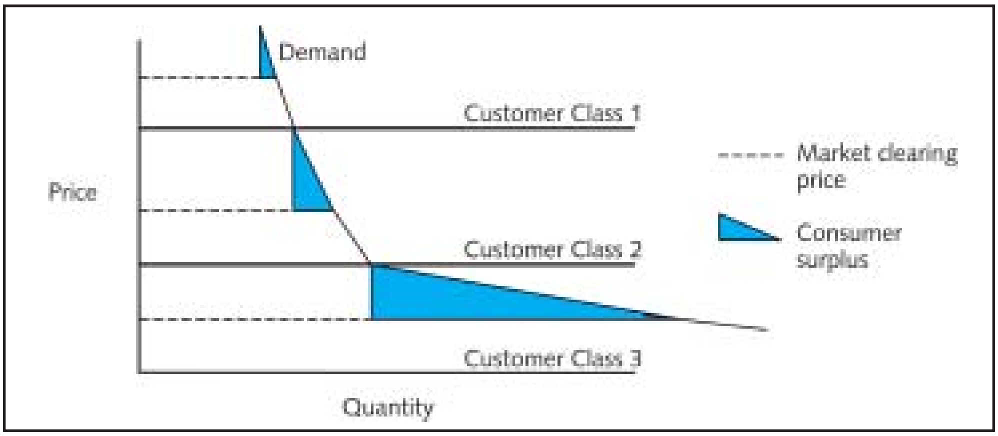

TABLE 12.4: *Airline Prices for an Identical Seat*

City Pair	Airline	Ticket Prices Full Coach	Advance Purchase 7 Days	Cheapest Fare
Washington, D.C.–Nashville	USAir	$ 811	$761	$251
Newark–Salt Lake City	Continental	$1,317	$571	$257
Dallas–Cleveland	American	$ 639	$471	$304

Source: Airline Web sites, March 2004.

Several difficulties can arise in creating these customer segments. Foremost, price discrimination, per se, is illegal in the United States. A business cannot simply tell some customers they must pay more for a good or service than others, just because the business knows that a customer will pay. Some types of customer segmentation are illegal in the United States, as well, such as segmenting customers according to race or ethnic origin and charging differential prices. This regulation does not apply in some other countries, however, and "ethnic pricing" of airfares is not an uncommon practice (Mitchener, 1997).

Consequently, it is up to the imagination of the marketing department to determine a legal and enforceable method for segmenting markets. For example, airlines want to charge higher prices to price-insensitive business customers and lower prices to price-sensitive vacationers. However, these markets can only be attacked indirectly, by segmenting pricing on how far in advance one makes a reservation.

The yield management issues of overbooking and capacity allocation lend themselves well to numerical analysis once pricing is set, but determining the best overall combination of prices, capacity allocation, and overbooking is difficult. In practice, marketing usually sets prices in coordination with company policy and competitive response, and the operations area sets capacity allocation numbers after pricing is determined. To show why pricing is so tricky in the yield management environment, consider the following example.

EXAMPLE 12.3: *Pricing and Capacity Allocation*

Consider a one-time event in which, miraculously, perfect economic information is provided on two groups of customers: Premium and Discount. For the Premium customers, the more that is charged, the fewer will attend, according to Table 12.5.

Consider three scenarios: (1) a service facility not facing a capacity constraint; (2) a capacity of only 100 customers, and Discount customers pay $50; and (3) again a capacity of 100 customers, but Discount customers pay $75. In scenario (1), the best price to charge premium customers is $90, resulting in a total of $10,800 in revenue, based on traditional economic logic. In a situation in which demand responds disproportionately to price cuts, prices should be low (high price elasticity, in economic-speak). Capacity restrictions and different customer classes, however, turn the traditional economics of price elasticity on its head. In scenario (2), high price elasticity still remains with the Premium customers, but one cannot take advantage of it, because the sheer number of potential customers overwhelms capacity. The best price to charge Premiums is $100, resulting in $10,000 revenue from dedicating the facility just to Premium customers and ignoring the Discount market. However, if the Discounts pay a bit more, as in scenario (3), the best price to charge Premiums is $110, resulting in more revenue by adding Discount customers.

TABLE 12.5: *Pricing and Capacity Allocation Example*

Premium Demand Information:			
Possible unit prices	$ 100	110	90
Associated demand	100	80	120
Scenario 1:			
Unlimited capacity, only Premium customers*			
Premium price	$ 100	110	90
Premium demand	100	80	120
Total revenue	$10,000	8,800	10,800
Scenario 2:			
Capacity of 100, Discount class unlimited demand at $50*			
Premium price	$ 100	110	90
Premium demand	100	80	100
Premium revenue	$10,000	8,800	9,000
Discount revenue	$ 0	1,000	0
Total revenue	$10,000	9,800	9,000
Scenario 3:			
Capacity of 100, Discount class unlimited demand at $75*			
Premium price	$ 100	110	90
Premium demand	100	80	100
Premium revenue	$10,000	8,800	9,000
Discount revenue	$ 0	1,500	0
Total revenue	$10,000	10,300	9,000

*Optimal solution in blue.

The overall lesson for pricing in yield management environments: Traditional reasoning regarding price elasticity does not apply, and pricing depends on the relative demand/capacity relationships that must be judged on an individual case-by-case basis.

Implementation Issues

Although yield management systems enjoy success in a number of different industries, some practical problems must be addressed in both customer relations and employee relations.

Alienating Customers

The chapter began with some examples of yield management practices that can alienate customers. From the customers' viewpoint, many of the rules for getting into a specific fare class seem ridiculous. The idea that they are told a service is full to capacity when they then find out it is not is unbelievable to many consumers. Further, pricing issues can create significant ill will among customers. The general public seems to accept the idea that one person may pay a different rate for an airline seat than another, although many people still get quite angry when they find out. That feeling, however, has not spread to other industries. If someone overhears the person in front of her in line getting a better room rate at a hotel, she will often *demand* that rate, even though the rate of her own customer class would have been perfectly fine with her before she knew about the better rate. A key in many industries is—to be blunt—to keep customers ignorant of the rates available to other customer classes.

Customer Class "Cheating"

If customers become knowledgeable about a system, they can manipulate it to their own benefit. For example, as stated previously in this chapter, business-oriented hotels are more willing to bargain with customers who are providing revenue on otherwise empty weekend nights. Consequently, one tactic used by some customers is to make a hotel reservation for three nights, Thursday through Saturday, then stay only Thursday night, their true original intention. Another tactic works in reverse: If a hotel is booked solid during weekdays, make a reservation for only Saturday night, then refuse to leave the hotel for several days. In many states, laws preclude evicting such guests. Even if the law allows, hoteliers find the process of evicting a paying customer both embarrassing and poor public relations, because it often must be done in front of other guests.

The complexity of airline customer class regulations created a veritable cottage industry of customer cheating. A number of firms specialize in the illegal (by airline rules) buying and selling of frequent flyer miles, and a few travel agents find sophisticated methods to avoid airline regulations as an added service to their customers. The airlines invest substantial sums in information systems to keep customers obeying the rules. For example, customers used to be able to avoid the expense of an expensive, mid-week round trip by purchasing two overlapping round-trip tickets with Saturday night stopovers and throwing away the other half of both tickets. Many airlines now, however, have the capability of detecting such abandoned tickets and require the customer to pay the appropriate price or lose the privilege of flying with that airline.

The solution to this type of problem usually requires both employee vigilance and investments in software.

Employee Empowerment

Strategic decisions must be made concerning how much power is given to employees to make decisions. If a system is seen as reducing employee discretion, employees can often make up for this perceived indignity by sabotaging the system. In the hotel industry, for example, a significant degree of responsibility for assigning the proper customer class and room rate resides in the clerk who sits at the front desk. If a system lacks enough structure, clerks may whimsically decide whether a customer gets a higher or lower rate depending on their mood rather than the customer's true class. However, a system with too much structure can lead to a loss of bargaining power by clerks, which can result in needlessly empty rooms. Consequently, it is important for anyone with such responsibility to understand the overall managerial goals of such a system.

Cost and Implementation Time

Fully implementing a yield management system for a large firm is not a simple matter. No off-the-shelf software can be purchased at the local mall to address yield management issues. The leading companies providing yield management software are Sabre, PROS Strategic Solutions, Manugistics, VeriTec Solutions, Opus 2 Revenue Technologies Inc., and VeriProfit, Inc. Implementation costs for such software can run several million dollars and it can take from several months to two years to implement the software. Further, once a system is running, training and system updating are also costly.

Summary

Yield management systems are used in a wide variety of industries that have limited capacity. Such systems can potentially make a significant difference in profitability for firms that use them well.

Three basic components of a yield management system are overbooking, capacity allocation, and pricing. For overbooking and capacity allocation, numerical methods presented in the chapter can help to solve those problems. However, a practical answer to the pricing problem still remains out of reach. Examples were used to show how the environment of a yield management system, with customer classes and capacity restrictions, can turn normal pricing decisions around.

Finally, a yield management system is not just a computer program. It is a system that must be implemented with flesh-and-blood employees and customers. These human elements of the system must be attended to carefully because poor employee or customer education as to the limits and nature of such a system may lead to its circumvention in a variety of ways.

Review Questions

1. A specific problem of yield management techniques noted in the chapter demonstrated that an optimal solution to a particular problem involving a 100-seat plane gave an average of 13 seats to discounters, 75 seats to premium passengers, and an average of 12 seats were unoccupied. Explain in words, not numbers, how it is possible for the best solution to leave seats unoccupied on average.
2. If the average number of no-shows are, for example, 12, why isn't the best overbooking number necessarily 12?
3. What are nested capacity allocation methods and why are they used?
4. Why can't the standard formula for overbooking be used in all cases?

Problems

12.1. The Greybeard Busing Company is assessing its overbooking policy for the Miami–Fort Myers run. The number of customers who don't show up after reserving a ticket is uniformly distributed from 1 to 10 (10% chance of 1 no-show, 10% chance of 2, etc.). Tickets cost $25, and if the particular bus run is full, a passenger with a reservation is given passage on a rival bus line, a cost of $60. What should Greybeard's overbooking policy be?

12.2. It's normal for an airline to overbook a flight by 20% or so, but Amtrak will overbook its long-distance trains only by 5% to 10%, even though they have even more no-shows than the airlines do. What reasons would Amtrak have for overbooking so little?

12.3. The Kaluauluauhala Resort rents weekly condominiums by the Hawaiian shore, nearly all to families from Japan and the U.S. mainland. It is the only resort on the west side of the island. These condos rent for $1,500 per week, and guests typically spend another $500 (in terms of profit for the firm) during their stay. A guest who doesn't show up is subject to a cancellation fee of $750. If no condo is available for a guest with a reservation, he has to be transported either to the other side of the island or to another island to a similar resort and Kaluauluauhala picks up half the cost of the other resort (about the same as Kaluauluauhala's price, so the customer receives a "half-price" vacation) and the average $200 transportation charge.

No-shows by customers average three each week, evenly distributed between 0 and 6. What should the resort's overbooking policy be?

12.4. Consultant Air focuses business on high-paying McKinsey consultants, but to fill planes, it also carries the general public. It is flying a 100-seat jet from Atlanta to San Francisco and the consultant demand is normally distributed with mean 60 and standard deviation 10. Consultants pay $1,200 per ticket, the general public pays $150. Of course, the general public must book a flight several weeks ahead of time, while the consultants book flights at the last minute. How many seats should be opened up for the general public? On average, assuming 0 no-shows, how full will the plane be under this plan?

12.5 In addition to the customer classes in the previous problem, Consultant Air added a third class, "normal business people," who pay $500 per ticket. Demand from these customers is normally distributed with mean 20 and standard deviation 5. How does this class of customers change the solution to the previous problem?

12.6. Assume the curves shown in Figure 12.7 represent only the highest, lowest, and average demand for full-fare flights. Other curves are also possible. You are now four weeks before the flight and have received 30 reservations. Assuming a 150-seat plane, what should your capacity allocation policy be?

12.7. The curves for problem 12.6 really represent averages rather than exact relationships. If the final boarding numbers represented by the curve are an average result, where the actual boarding passengers are normally distributed about that average with a standard deviation of 10, what should your nested capacity allocation policy be? Assume prices of $1,000 per seat for premium, $400 for discount.

FIGURE 12.7: *Dynamic Capacity Allocation*

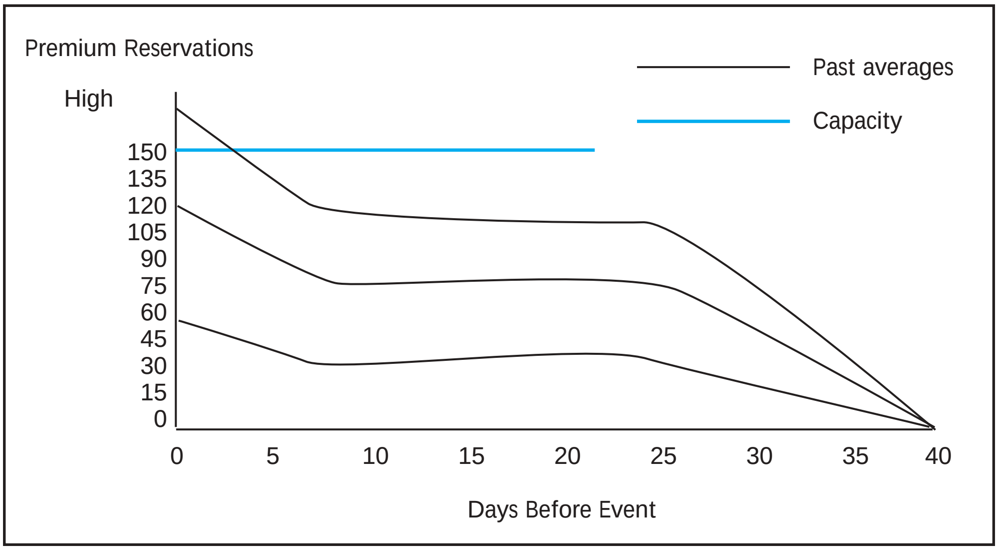

12.8. Rene was adamant. "If we have to tell customers we overbooked and there's no seating available, they are going to be extremely angry. I'll bet it costs us $150 per couple in goodwill each time." Tom replied, "No way, it doesn't cost us a penny. We're so popular now it doesn't matter if one potential customer doesn't come back. What matters is getting the most profit out of each seating. We prep everything ahead of time, so it costs us $50 per couple for food whether they show up or not, and we charge $150 per couple. At that rate of profit I don't want to see a single empty table."

The restaurant Chez Rogat was run by two Owen grads. They had two seatings each night, 6:00 P.M. and 8:30 P.M., for their 30-table restaurant. They prepared the same meal for everyone at each seating. Due to their popularity, they always had more reservations than they could handle and reservations were made far in advance. They had few no-shows, but did have some. They had about an equal chance of either 0, 1, 2, or 3 couples as no-shows at any given seating.

Given their costs, calculate an overbooking policy for Chez Rogat. Would you recommend any other capacity management strategy?

Selected Bibliography

Belobaba, P. 1989. Application of a Probabilistic Decision Model to Airline Seat Inventory Control. *Operations Research*, 37(2), 183–197.

Cook, T. 1998. SABRE Soars. *OR/MS Today* (June), 26–31.

Crandall, R., and R. Markland. 1996. Demand Management: Today's Challenge for Service Industries. *Production and Operations Management*, 5(2), 106–120.

Cross, R. 1997. *Revenue Management*. Broadway Books, NY.

Cross, R. 2004. February 10 presentation at Emory University.

Mitchener, B. 1997. "Ethnic Pricing" Means Unfair Air Fares. *The Wall Street Journal* (December 5), B1.

Smith, B., J. Leimkuhler, and R. Darrow. 1992. Yield Management at American Airlines, *Interfaces*, 22(1), 8–31.

CASE STUDY

Yield Management at MotherLand Air

MotherLand Air is an airline dedicated to the betterment of the proletariat, with original operating principles based upon the strong political beliefs of its founder, Al Niemi. Al started this airline for the selfless purpose of exposing Dallasites to the glories of communist ideals. His own vacation in North Korea convinced him that the communist economic system held ultimate truths, and he reasoned that his colleagues would be likewise convinced if only they would visit the motherland. Consequently, he leased a Boeing 747 and initiated direct service from Dallas to Pyongyang three days a week.

Al was philosophically against such capitalist notions as differential pricing for customers. He believed that all comrades were equal and should pay the same fare. Consequently, MotherLand Air was started as a one-price airline offering round-trip service for $400. Unfortunately, his 400-seat 747 was flying at far less than capacity seating on most trips and MotherLand was highly unprofitable.

Al was stunned that more travelers did not desire the icy North Korean air in their lungs, but he knew he had to adapt or shut down. He began to see wisdom in the old saying, "We are all equal, but some are more equal than others." Accordingly, he tried to salvage MotherLand by utilizing differential pricing. He segmented his customers into three different customer categories—"capitalist class" business flyers (who paid full coach fare), "comrade class" discount tickets, and "party member class" deep discount tickets—based on length of stay and flying days and time of year. He charged $1,000 for business flyers, $400 for discount tickets, and $100 for deep discount tickets.

This move improved his revenue significantly. He typically sold from 30 to 50 full-fare seats (Table 12.6), 70 to 100 discount seats (Table 12.7), and 100 to 150 deep discount seats on each trip.

TABLE 12.6: *Full Coach Fare History: Net Reservation Activity*

		Weeks From Event							Months From Event				
Obs.	Boarded	1	2	3	4	5	6	7	2	3	4	5	6
1	34	53	51	37	28	23	16	15	13	8	9	6	5
2	26	40	40	29	18	16	14	12	10	6	7	5	5
3	27	42	39	30	22	15	15	13	10	7	6	6	4
4	36	43	40	28	21	17	11	10	11	6	6	5	4
•													
•													
•													
39	33	42	37	34	18	17	13	11	10	7	7	5	4
40	40	56	51	38	29	23	15	14	14	12	8	7	7

CASE STUDY

Access your Student CD now for Tables 12.6 and 12.7 as Excel spreadsheets.

TABLE 12.7: *Discount Fare History: Net Reservation Activity*

		Weeks From Event							Months From Event				
Obs.	**Boarded**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**2**	**3**	**4**	**5**	**6**
1	87	88	90	92	90	84	76	71	59	47	36	24	14
2	81	85	87	88	87	81	75	67	60	46	34	21	12
3	65	70	71	73	70	69	67	62	53	41	29	22	16
4	57	57	58	59	62	58	55	47	38	31	23	16	11
•													
•													
•													
39	100	101	103	105	104	102	95	85	71	57	44	32	21
40	105	111	113	115	114	110	104	93	79	65	50	32	19

Unfortunately, even with his new pricing scheme he was still not able to attain profitability, but a solution became clear. Since his 747 was never full, he negotiated a change in his lease for a much cheaper and more fuel-efficient 100-seat plane that would allow him to become profitable.

At this point, however, he needs assistance. He could fill his planes with deep discount and discount flyers, but it would be unprofitable to do so. Further, he had no experience with overbooking policies and was unsure whether he wanted to pursue one. Consequently, Al needs expert service operations consulting to maximize revenue.

Seating Allocation

Seating allocation must come in a "nested" form. For example, in Table 12.8 no deep discount, up to 50 discount and deep discount, and up to a combined 120 full coach, discount and deep discount seats are set aside six months prior to departure. Implicitly, for a 100-seat plane this allocation contains 20 overbooked seats.

Demand history from 40 flights on the larger planes is contained on Tables 12.6 and 12.7 and is available on the text CD. For example, in the Full Coach Fare History, Observation 1 indicates that 5 reservations were obtained 6 months before the flight, 6 reservations obtained 5 months before the flight, and so on, and of the 53 reservations in the system the week before the flight, 34 actually boarded the plane. Demand history for the deep discount tickets is not necessary because 100% of them are taken as soon as they are offered.

The demand history given can be used to determine threshold curves for each class of passenger. The purpose of developing threshold curves is to predict the eventual number of reservations just prior to boarding.

The costs involved in overbooking depend on the number of people not seated for the flight. The revenue from the ticket sales is subtracted (for simplicity, assume revenue of $1,000 for each person not seated), and Al estimates the per passenger

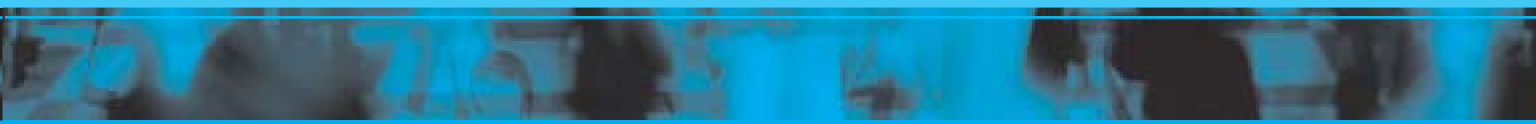

CASE STUDY

TABLE 12.8: *Game Tracking Sheet*

Flight Group Name:

	Weeks Away from Takeoff											
	24	20	16	12	8	7	6	5	4	3	2	1
Price												
Full	1,000											
Discount	400											
Deep Discount	100											
Price Effects on Reservations												
Full												
Discount												
Policy												
Full	120											
Discount	50											
Deep Discount	0											
Booking												
Revenues												
Full												
Discount												
Deep Discount												
Revenue												
Overbooking Penalty												
Profit												

penalty cost of leaving passengers stranded as $200 times the number stranded squared. So leaving one person stranded costs $200, two people costs $800, three people costs $1,800, and so on.

Pricing

Currently, ticket prices start at $1,000, $400, and $100 for the three fare classes. Prices can be raised by 10% to $1,100 or $440, lowering demand by 15%, or prices can be lowered by 10% to $900 or $360, increasing demand by 15%. Prices on deep discount tickets cannot be changed.

Assignment

As a consultant, you advise Al Niemi concerning overbooking, seating allocation, and pricing for the three classes of passengers. Your group must indicate a policy prior to receiving any reservations. Once all policies are in, the reservation information for the six-month-out time frame will be given. At this point, you may alter your policies. Iterations of receiving reservation information and altering policies proceeds until plane departure, as depicted in Table 12.8.