Evaluating Capacity- and Demand-Management Decisions at a Ski Resort

This model reveals some effective and not-so-effective ways to manipulate customers' behavior to minimize their waiting times in service queues.

BY MADELEINE E. PULLMAN AND GARY M. THOMPSON

Capacity planning is a dominant issue when developing hospitality-operations strategies. From relatively small-scale settings such as restaurants to massive venues such as cruise ships and theme parks, managers face the same questions, such as: How can the service create the most revenue from a limited space and still keep the guests satisfied?; Should the business promote off-peak use at discount prices, add more capacity (at the risk of making guests feel crowded), or target marketing campaigns to those customer groups that might take advantage of underused facilities?; and Should the firm implement automated reservation technology to control queues for a ride or attraction?

Depending on the characteristics of the hospitality operation, those capacity decisions determine other measures of service management such as productivity, growth, change, and competition.[1] The well-known core problem for the service industry is that demand for services must be met as it arises, because it cannot be inventoried. Demand variability creates alternating periods of idle service workers or facilities and consumer waits.[2] Therefore, management must trade off the cost of idle resources versus the potential cost of customer dissatisfaction with long waits. Dissatisfied customers hurt the long-term profits and success of service firms for the following reasons: (1) failure of the customer to return for future business; (2) reduction in the customer's frequency of

© 2002, CORNELL UNIVERSITY

[1] J.A. Fitzsimmons and M.J. Fitzsimmons, *Service Management for Competitive Advantage* (New York, NY, McGraw-Hill, 1994).

[2] M.J. Maggard, "Determining Electronic Point-of-sale Cash Register Requirements," *Journal of Retailing*, Vol. 57, No. 2 (1981), pp. 64–86.

visits; and (3) negative word-of-mouth advertising.[3] Thus, managers must consider the long-term effects of their capacity strategies. In this paper we examine the issue of capacity planning in a service network, namely, at a popular U.S. ski resort.

Service networks are businesses that offer multiple services and activities at one site. Hospitality examples include theme parks, casinos, cruise ships, airlines' multimedia entertainment systems, and exercise facilities. Service networks face more-complex capacity-management decisions than do services with a single service offering. This complexity occurs due to intra-activity demand variation (e.g., the popularity of different activities or restaurants on a cruise ship varies from day to day), the number of activities and servers available at any given time, and physical constraints such as space limitations. In a service network, customers pay one fee to enter the site, pay per activity visited on the site, or both. Once the customer enters the site, the individual either chooses or is directed (e.g., as part of a package tour) to participate in any number of activities within the system. Customers randomly arrive at the activities, and each activity has a different duration for each customer. As a result, lengthy waits can occur at each activity depending on the service times. This problem is apparent at ski resorts, where patrons must queue up to wait for the various lifts (not to mention the resort's various food-service outlets).

> Simulation provides analysts with a method for modeling various "what if" scenarios in a simulated environment before spending real money to implement untested ideas on real customers.

The goal, of course, is to maximize the number of customers entering the site (e.g., amusement park) or paying for facilities within the site. On the other hand, customers would like to minimize their waiting time during any encounter (e.g., a ride on the tallest roller-coaster). As Lovelock aptly put it, "nobody wants to be at Disney World on a record-breaking day for ticket sales."[4] Service networks offer distinctive opportunities for matching supply to demand within the system. Assuming that managers have some influence over external demand through their marketing efforts, the operations area is typically responsible for reallocating demand within the system and matching supply to demand on a real-time basis (e.g., by moving patrons from one attraction to another). Although it is possible to effectively manage capacity through a single strategy or perspective, an optimal strategy would require an integrated set of strategies representing both demand and supply perspectives.

Assuming that (a) hospitality firms want to maximize capacity use for a given level of customer satisfaction in a constrained-capacity environment, (b) customer satisfaction is directly related to overall waiting time in the service system, and (c) management has the capability to manage certain aspects of supply and demand, the research question driving our investigation is:

> What is the optimal combination of supply- and demand-management strategies that provide for maximum capacity in the service system while operating within a specified peak waiting-time limit for customers?

Our primary objective in this study was to evaluate and recommend the appropriate alternatives for capacity management. By examining all possible combinations of alternatives, we determined the capacity strategy that maximizes the percentage of customers experiencing a wait time of less than ten minutes. Though our investigation is based on a ski resort, our findings have implications for any hospitality firm considering some combination of capacity expansion or demand-management strategies. We use a service-standard metric—the proportion of customers who experience a wait of 10 minutes or less—to evaluate the effect of the capacity and demand-management strategies. Our informal discussions with skiers suggest that a 10-minute service standard is reasonable.

[3] M.M. Davis, "How Long Should a Customer Wait for Service?," *Decision Sciences*, Vol. 22 (1991), pp. 421–434.

[4] C. Lovelock, *Managing Services: Marketing, Operations, and Human Resources*, second edition (Englewood Cliffs, NJ: Prentice-Hall, 1992).

EXHIBIT 1

The resort's ski-lift layout

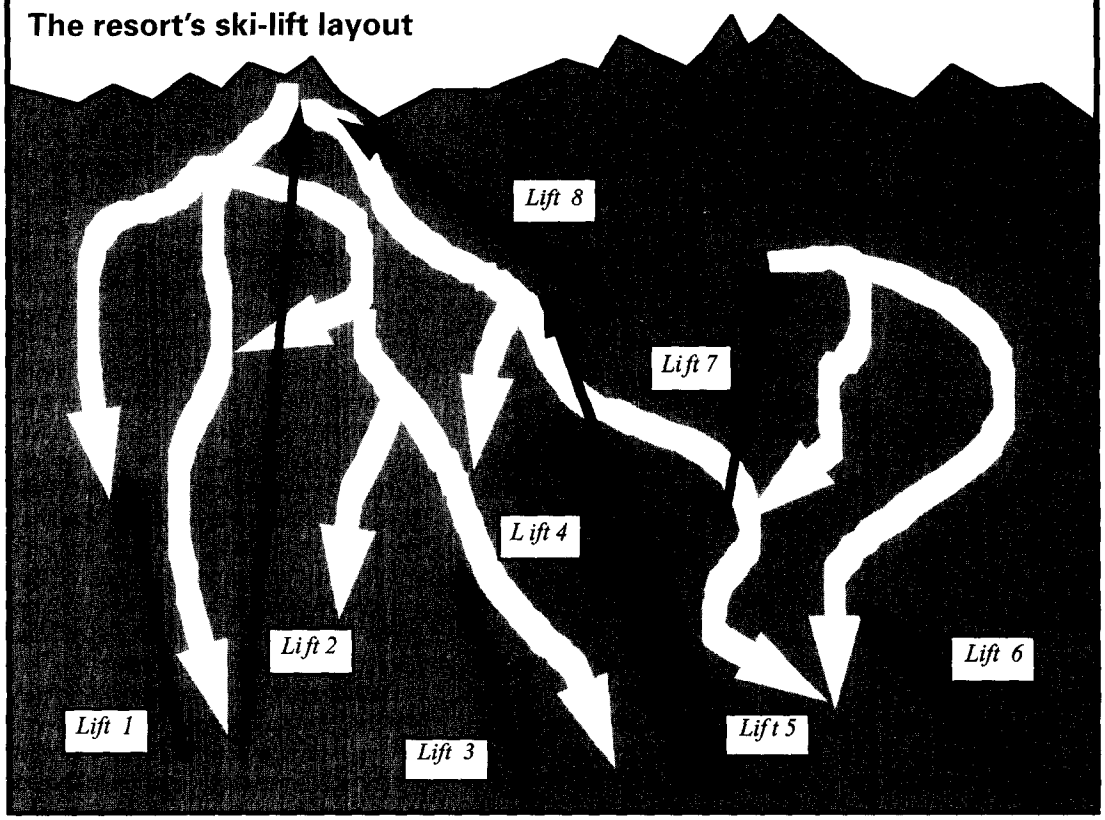

Simulation Modeling

For studying hospitality businesses under many different capacity-management scenarios, simulation modeling is a useful tool. Simulation provides the analyst with a method for modeling actual customer demand and server-variability and -shift patterns, evaluating the effects of different strategies in a simulated environment before implementing the idea with real people, and generating useful performance measures and reports such as employee-use rates, wait times, and machinery-breakdown times. Many of the current software packages have graphic-output capabilities for modeling typical hospitality settings such as restaurants, airports, cruise ships, and theme parks. For example, the analyst can watch how a restaurant might perform throughout an evening in both front- and back-of-house areas when a certain promotion increases customer traffic by 10 percent or moves 25 percent of orders from the sauté station to the grill. Simulation software ranges from a simple and inexpensive package, Roller Coaster Tycoon (a theme-park simulation game), to the highly customizable ServiceModel, which incorporates graphics and metrics for most hospitality applications.

A typical simulation exercise involves collecting data on existing system characteristics, determining which variables will change (i.e., number of seats, speed of service, types of customers, and number of employees on duty), setting up the initial simulation and confirming that it matches existing performance, running different scenarios, and evaluating the results. In the next section, we show how this approach is applied to our ski-resort problem.

Data Collection at the Ski Resort

The data for this study were collected from a ski resort in Utah. Ski resorts exhibit all attributes of a service network. They serve various customer classes (beginner through expert);[5] use a system

[5] Customer classes are groups of customers having similar characteristics.

Demand- and Capacity-data Collection

To generate our forecasts of daily demand, the resort provided us with the last ten years of historic customer and ski-condition data, which included daily customer tallies, snow depth, and new snowfall, by month and day type (weekend, weekday, or holiday). Additionally, members of the resort's management team described feasible capacity improvements and the technology that could advance those improvements. The possible improvements included: new high-speed chair lifts with double the seating capacity of existing chair lifts; and expansion into new terrain with a new lift and new runs. Managers had proposed a restricted weekday ticket, which they estimated would shift 10 percent of the weekend skiers equally among the weekdays. The managers also estimated that demand could potentially increase with technology improvements and trends in the industry, ranging from 5 to 20 percent from existing demand.

The resort's managers had several alternatives to shift demand within the ski-lift network. Using an information system, the resort could post electronic displays (signs) to announce the estimated waiting times at the various lifts. With this feature, managers estimated that 50 percent of the customers would shift to a lift with a shorter waiting time—depending on their skiing ability. Another option would be to introduce promotions to attract more customers from the less-expert classes. (The existing resort is patronized predominantly by customers from high-ability levels.) More nonexpert skiers could potentially affect the demand distribution for the various lifts. (Increasing overall customer volume will change the demand for the different lifts and will cause the waits to change. A simulation can predict what those changes will be.)

Existing service time. We collected service-time data by observing all of the lifts. An observer measured the cycle time for each chair and interviewed the lift attendant to determine the frequency of stoppages for the lift. This resort's ski-lift network consists of eight different chair lifts, each with a different seating capacity, lift speed, and ride duration (see Exhibit 1). All lifts suffer from occasional random stoppages, which increase proportionally with **(a)** increased demand on the facility, and **(b)** number of beginners using the facility. To model the probability of a lift's having to stop unexpectedly, we used the current wait time at the lift as a proxy for demand level. We assumed that the probability of a stop is a linear function of the current wait time and percentage of beginners in the lift line. When a stop occurs, we added the downtime to the lift time for all customers on that lift.

Daily arrival pattern for customers. The resort has two major rush periods: the first, from 8:00 AM until 9:30 AM; and the second, from 11:30 AM until 1:00 PM. We collected baseline daily arrival-pattern data on three sampling days, two weekends, and one weekday. To determine the arrival pattern, we counted the number of customers walking through the resort entrance for five-minute intervals every 15 minutes, starting 15 minutes before the opening of the ticket-sales window. We then compared the arrival patterns to the total number of skiers for the day and day type (weekend or weekday). Based on the observed arrival patterns and discussions with managers, we simulated the variability in the percentage of afternoon (after-11:30 AM) skiers and the percentage of skiers falling in each arrival block (8:00 AM–9:30 AM, 9:30 AM–11:30 AM, 11:30 AM–1:00 PM, 1:00 PM–2:30 PM) to develop, for each day, a continuous, piecewise, time-dependent, linear empirical distribution of the time between skier arrivals at the resort.

We collected validation data for waiting statistics at each lift by observing an entire day for two sample days. An observer was stationed at a lift and recorded the number of people in line and the waiting time for the last person entering the line. The waiting time is the difference (in minutes and seconds) between when the person joins the queue and when the person enters or sits on the lift. Our observers collected those data at each lift every 30 minutes for the entire day (i.e., from **8**:00 AM until **5**:00 PM). The resort provided data on the number of customers entering the system each day. The simulation results for the peak-wait data at each lift should correspond to the actual statistics gathered for a certain number of skiers entering the system.

Travel-time observations. Our observers recorded the travel time between service facilities on ten different days during the ski season. The observers averaged ten observations per day for a total of 100 observations. During each day of the study, skiers were observed on two of the eight possible lifts. An observer randomly selected a customer departing a lift and followed the customer until she arrived at the next lift. The observer noted the skier's ability (beginner to advanced), run choice, weather and terrain conditions, and travel time between facilities. On the days sampled, all runs and lifts were open. Additionally, a group of expert skiers provided us with information on the minimum times that it would take them to move between facilities.

Daily network-flow patterns. To determine the probability of customers' moving between activities as a function of their customer class, we administered a "skier diary" survey to skiers during their lunch break or after skiing (an example of which is presented in Exhibit 3, on page 31). The survey asked skiers to outline their previous choices of lifts and connecting runs for either the morning or the afternoon period. We summarized this information to develop an empirical frequency-distribution matrix for the existing resort. For the improvement scenarios, the probability of a customer's choosing an upgraded lift equals the original probability. To model the results of expanded capacity, we set the probability of a skier's choosing a new lift equal to that of a comparable lift (i.e., by location and type of terrain), then reweighted all the probabilities to sum to one.

The instrument also asked skiers to provide an estimation of **(1)** their skiing ability, **(2)** their arrival, departure, break, and lunch times, and **(3)** other demographic information. We used this additional information in the simulation to generate normal distributions for each customer's ski time until lunch, lunch duration and location, and ski duration until departure time.—*M.E.P. and G.M.T.*

of activities (ski lifts) where queuing occurs; and can predict movement between queues based on customer ability (e.g., beginners tend to avoid difficult terrain while experts head for the challenging runs). Exhibit 1 (on page 27) shows an illustration of the ski resort's lift configuration. The data-collection phase involved several steps: (1) estimation of daily demand, (2) determination of feasible demand-shifting options and capacity improvements, (3) determination of existing service time for each activity, (4) daily arrival rates for customers at each lift during each time period, (5) daily network flow patterns for different customer classes, (6) time for travel between lifts as a function of customer class, and (7) validation of the existing configuration simulation. We describe each of those steps in the accompanying box (at left).

The Ski-resort Simulation Experiment

We selected the methodology of simulation since it is ideally suited for representing the variability and interrelationships that exist in service networks.[6] Our simulation experiment has six main steps. First, we set up the configuration of the simulation to match the resort's existing configuration. Second, we generated the daily conditions and a customer-demand pattern for the day. Third, we ran the simulation for a simulated day and collected queuing statistics. Fourth, we ran the daily simulation for ten hypothetical, simulated years. Fifth, we compared the configuration waiting-line results with actual data. (When necessary, we adjusted the existing configuration model assumptions until our simulation results were comparable to our actual resort data.) Finally, we changed the experimental factor level and repeated steps 2 through 4 for each configuration. We describe those steps in more detail below.

Network configuration and strategy routine. The network has many possible configurations of capacity- and demand-management alternatives, as shown in Exhibit 2. The baseline configuration is the existing network and demand patterns. We classify strategies as (1) managerially controllable and (2) environmental, which are those indirectly influenced by managers or being outside of management's control. The managerially controllable variables include lift-capacity upgrades, terrain expansion, demand smoothing via ticket-pricing alternatives, and queue-information efforts (e.g., signs for skiers about lift conditions). Environmental variables consist of demand growth and customer-mix (customer class) variations. Capacity is increased to three possible levels using lift upgrades or one additional level of terrain expansion. Information effort (i.e., the use of signs to announce lift conditions) has three levels: (1) the no-information scenario, assuming customers ignore previous wait experience when selecting their next lift; (2) the personal-wait-knowledge scenario, assuming customers do not immediately repeat the pre-

EXHIBIT 2

Experimental factors

Factor	Number of levels	Unit of measure	Levels
Capacity upgrades: *Improved lifts*	4	Uphill capacity: customers per cycle	1) No change 2) 1 high-speed quad 3) 2 high-speed quad 4) 3 high-speed quad
Capacity expansion: *Additional terrain*	2	Uphill capacity: customers per cycle	1) No change 2) 1 high-speed quad with new terrain
Information use	3	Use of prior and current queue information	1) Customers return to any lift 2) Customers do not return to the lift if the wait > average 3) Half of the customers move to shorter queues
Industry growth	3	Proportion of existing daily demand	1) No change 2) 5% growth 3) 20% growth
Inter-daily demand smoothing	2	Proportion of weekend demand moved	1) No change 2) 10 percent of weekend to weekdays
Class variation	2	Proportion of customer classes	1) Existing mix 2) Increase beginners and intermediates

[6] For an introduction to the use of simulation in hospitality research, see: G.M. Thompson and R. Verma, "Computer Simulation in Hospitality Teaching, Practice, and Research," *Cornell Hotel and Restaurant Administration Quarterly* (scheduled to be published in April 2003).

vious lift if their wait was longer than their average wait; and (3) the queue-information scenario, where customers have a 50-percent chance of using queue information from signs to find the shortest nearby, skill-appropriate lift line. All of those managerial decisions occur under three levels of industry growth (i.e., none, 5 percent, and 20 percent from existing demand, assuming that each test year has the same growth), demand smoothing with two different alternatives of ticket pricing, and two levels of customer-class mix (existing mix and desired mix). For a full factorial experimental design, 288 different scenarios exist ($4 \times 2 \times 3 \times 3 \times 2 \times 2$).

Daily conditions assignment routine. We used multiple regression to forecast the number of customers arriving on a given day. The regression equation is based on 250 days of historic data, randomly selected from the past ten years. While the original demand model included the previous day's and the current day's inches of new snow, we found that those variables were not statistically significant. Therefore, the daily demand is a function of just these statistically significant variables: existing snowpack, month, and day type.[7]

With an adjusted R^2 of .368, the explained variance is relatively low, indicating a relatively high amount of intrinsic variation in daily demand.

The simulation logic that determines the snow conditions and number of skiers who arrive on a given day is as follows. Starting with the third Wednesday in November, the program randomly generates an existing snow pack and day's weather based on a normal distribution of historic conditions for that day. The program then determines whether the conditions are appropriate to open the resort (i.e., the snow pack must be above a predetermined minimum). If the resort cannot open, the preceding steps continue until the snow pack builds up to an appropriate depth. The day number incrementally increases with each repetition. Snow pack increases with new snowfall or, if there's no new snow, decreases by 1.5 percent each day due to settling.

When the snow pack is adequate, the program determines whether the day type is appropriate to open for the season. The opening day must be a holiday or weekend day. Additionally, this stage of the routine determines when the resort should be closed at the end of the season based on snow pack. The routine also reduces demand to 30 percent of the original value if too much snow falls (e.g., greater than 24 inches of snowfall in one day). If the resort is open, the program determines how many skiers will arrive during the day—specifically, the number of skiers that arrive in the morning and those that arrive in the afternoon. The program runs the daily simulation and collects queue statistics for the complete day, after which it resets its counters. Again, the program updates the day number, month, and day type for a new day—along with new weather conditions. As the season progresses, the program checks on the day number and snow-pack conditions. If the snow pack falls below a certain minimum or the day number exceeds 155 days, the resort is closed for the season.

The daily conditions routinely generated 100 hypothetical seasons of daily demand patterns. Of those 100 patterns, we chose ten test years to replicate the entire model. Two of those years represent low-demand years, two represent high-demand years, and six represent average years. The simulation program runs each scenario for each day of the ten test years, or approximately 1,550 sample days.

Finally, the simulation program randomly assigns an arrival time for each day's customers. Each day of each test year has a specific random number stream so that each capacity–demand scenario has identical arrival patterns and demand for a particular day. The program determines the departure time, lunch time, and ski-ability level for each customer using empirical frequency distributions derived from the "skier diary" survey.

Simulations. The program simulates each hypothetical skier's activities for the day, as follows:

(1) The skier arrives and is assigned an ability level, lunch time, departure time, and departure location.

[7] Using multiple regression, we found that the number of daily skiers, c_d, to be given by $c_d = 1349.48 + 7.09 \times SP_d + 805.80 \times W_d + 661.63 \times H_d$, where SP_d is the cumulative snowpack, in inches, on day d; $W_d = 1$ if the day is a weekend, $W_d = 0$ otherwise; and $H_d = 2$ if day d is a major holiday, $H_d = 1$ if a minor holiday, $H_d = 0$ otherwise.

EXHIBIT 3

Sample skier

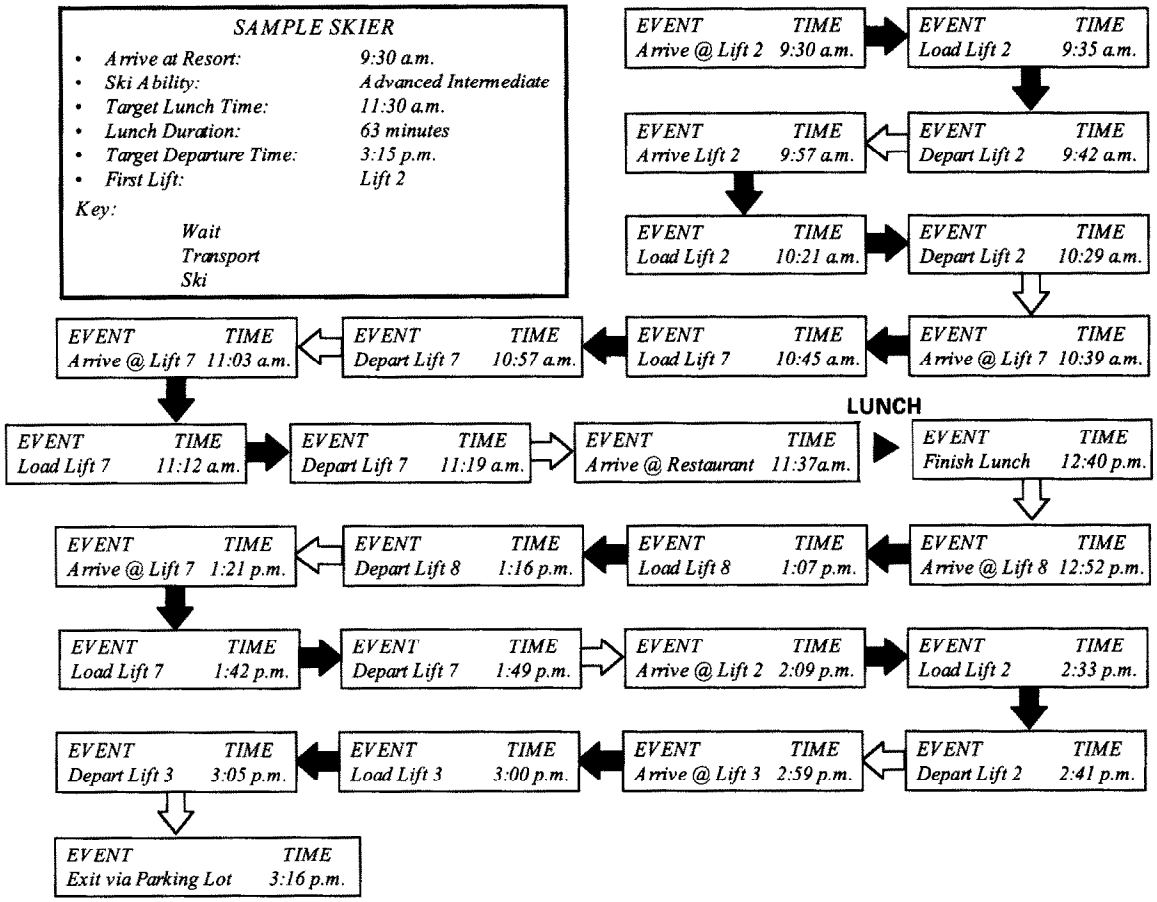

(2) The skier chooses a lift with a probability dependent on her ability level and current location.

(3) The skier enters the lift line, and her waiting-line statistics begin. If there is no wait, the skier enters the lift and is kept on the lift for the appropriate lift travel time. When the skier enters the lift, the program determines whether a stoppage occurs. If so, the stoppage time is added to all the lift's loaded customers' transport time. Skiers do not leave the stopped lift's waiting line.

(4) At the end of each lift, the program checks to see whether the skier has passed the desired lunch time or departure time, minus 15 minutes. If this occurs, the skier goes to the lunch location or departs. Otherwise, the model finds the skier's next lift, her travel time to that lift, and repeats step 3.

(5) The duration of lunch is based on the normal distribution of empirical data. After lunch, the skier goes to step 2.

(6) Beginning at 3:45 PM, certain lifts close. Remaining skiers randomly choose an available open lift until all lifts finally close at 5:00 PM. When all lifts close, they depart from their last lift choice.

Exhibit 3 shows an example of a typical skier's simulated routine. The simulation is run for the entire day, and waiting statistics are collected for

EXHIBIT 4

Service levels for minimizing the ski-lift wait time based on *current* customer mix

		No inter-day demand smoothing			Use of inter-day demand smoothing		
		No prior queue information	Personal queue knowledge	Resort queue information	No prior queue information	Personal queue knowledge	Resort queue information
Capacity changes		(a) (b) (c)	(a) (b) (c)	(a) (b) (c)	(a) (b) (c)	(a) (b) (c)	(a) (b) (c)
No new terrain	0 new lifts	.74, .71, .60*	.77, .73, .61	.87, .84, .71	.75, .72, .60	.79, .74, .63	.90, .86, .73
	1 new lift	.72, .70, .63	.86, .84, .76	.95, .92, .84	.73, .70, .63	.87, .85, .77	.97, .95, .87
	2 new lifts	.73, .70, .63	.89, .87, .81	.98, .97, .92	.74, .71, .64	.90, .88, .82	.99, .98, .94
	3 new lifts	.89, .86, .79	.90, .88, .83	.99, .98, .93	.90, .88, .80	.91, .89, .83	.99, .98, .95
Expanded terrain	0 new lifts	.75, .71, .61	.86 , .83, .76	.92, .89, .79	.77, .73, .62	.87, .84, .77	.94, .92, .82
	1 new lift	.74, .71, .64	.87, .85, .77	.96, .94, .89	.75, .72, .64	.88, .85, .78	.97, .96, .91
	2 new lifts	.75, .73, .65	.89 , .87, .81	.98, .97, .94	.76, .74, .65	.90, .88, .81	.99, .98, .95
	3 new lifts	.90, .88, .80	.90, .88, .82	.99, .98, .95	.91, .88, .81	.91, .89, .83	.99, .99, .96

* Values are the proportion of customers who experience a peak wait time of 10 minutes or less under this scenario. Each triplet represents, respectively, the results with **(a)** no growth, **(b)** 5-percent growth and **(c)** 20-percent growth.

all customers and for each activity. The daily routine is repeated for the entire season and the seasonal routine is repeated for 10 hypothetical seasons.

Validation. The nature of the model and the problem that it examines makes validation difficult. To collect a large number of day-long waiting-line statistics requires an observer at each of the eight lifts, all day for many sample days. Therefore, we focused our model validation efforts in two areas. First, we compared the results of the model with the opinions of resort managers and a group of frequent skiers who have skied at the resort at least 50 times in the past five years. Second, we compared the results of the model against actual observed waiting-line data collected at certain lifts during different demand days. The frequent ski customers and resort managers provided estimates of when and how long peak waits would be for different demand levels at the resort. All the individuals agreed that lifts 2 and 4 would experience the longest waits as demand increased. Those waits could be as long as 45 to 60 minutes on days with demand exceeding 4,500 skiers. Predicting when peak wait time would occur is generally not possible, because the waiting time depends on snow and weather conditions during the day. After running the model for the existing configuration, the season-pass customers and resort managers determined that the overall model wait-time results adequately represented reality.

Results

Exhibits 4 and 5 provide the service standards for minimal peak wait (e.g., the proportion of customers who experienced peak waits of less than 10 minutes).[8] Exhibit 4 shows the scenarios with the current customer-class mix, while Exhibit 5 provides the comparable scenarios with the managers' desired customer-class mix. Those exhibits show the capacity changes on the left column, and the smoothing option and the rate of queue-information use on the top rows. The service standard for the resort's existing configuration and current demand, the first row in Exhibit 4, indicates that 74 percent of the customers experience a peak wait of 10 minutes or less assuming that they do not use any previous knowledge of queue lengths for their lift choices.

[8] Our results are not sensitive to our selection of a 10-minute wait limit. For example, our findings would be similar had we selected a 15-minute wait limit (though the proportion of skiers experiencing this wait would be different).

EXHIBIT 5

Service levels for minimizing the ski-lift wait time based on *desired* customer mix

		No inter-day demand smoothing			Use of inter-day demand smoothing		
		No prior queue information	Personal queue knowledge	Resort queue information	No prior queue information	Personal queue knowledge	Resort queue information
Capacity changes		(a) (b) (c)	(a) (b) (c)	(a) (b) (c)	(a) (b) (c)	(a) (b) (c)	(a) (b) (c)
No new terrain	0 new lifts	.53, .50, .45*	.67, .65, .59	.73, .69, .55	.52, .50, .45	.67, .66, .59	.76, .71, .56
	1 new lift	.46, .45, .42	.69, .67, .61	.81, .77, .67	.46, .44, .41	.69, .67, .61	.83, .79, .68
	2 new lifts	.47, .45, .42	.70, .68, .63	.87, .84, .77	.47, .45, .41	.70, .67, .62	.89, .86, .77
	3 new lifts	.67, .64, .58	.70, .68, .63	.88, .86, .78	.67, .64, .57	.70, ,68, .63	.89, .87, .79
Expanded terrain	0 new lifts	.53, .50, .45	.67, .65, .60	.78, .75, .65	.52, .50, .45	.67, .65, .60	.79, .76, .66
	1 new lift	.47, .46, .42	.69, .67, .61	.82, .79, .71	.47, .46, .42	.69, .67, .61	.83, .80. .71
	2 new lifts	.49, .46, .42	.69, .68, .63	.87, .85, .78	.48, .46, .42	.70, .67, .62	.89, .86, .78
	3 new lifts	.67, .65, .59	.70, .68, .63	.88, .86, .79	.68, .64, .58	.70, .68, .63	.90, .87, .80

* Values are the proportion of customers who experience a peak wait time of 10 minutes or less under this scenario. Each triplet represents, respectively, the results with **(a)** no growth, **(b)** 5-percent growth and **(c)** 20-percent growth.

If all customers use knowledge of queue lengths and do not repeat a previous lift when the queue length is longer than average, then the service standard increases to 0.77. The baseline service standard for the existing resort is assumed to be 0.74 because there is no way of knowing the extent to which customers rely on their previous queue knowledge.

Capacity changes and queue information. Given the existing demand environment, management has the ability to adjust the two capacity variables and provide current queue information. Exhibit 4 indicates that when existing lifts are replaced with faster lifts with more seating capacity, as long as customers use no prior queue information, the service standards stay the same or actually decline until three lifts are upgraded. The lifts are replaced in the following order: lift 4, lift 1, and lift 8. When lift 4—a major lift for access to the upper mountain—is replaced with a faster and higher-capacity lift, wait times increase a small amount at the four other lifts accessible from lift 4. This increase occurs because the skier-arrival rate at those lifts increases because skiers are departing lift 4 faster than before. Similarly, when lift 1 is replaced, waiting time increases at two other lifts. After replacing lift 8, the system has finally reached a balanced state, and waiting times decline in almost all other lift locations.

After adding three new lifts, the service standard jumps from .74 to .89 with no new terrain and from .75 to .90 with expanded terrain. Terrain expansion alone provides not more than a 2-percentage-point increase in service standard regardless of the number of upgraded lifts, all other variables held constant. Thus, if one assumes that customers do not use prior queue information, management must replace three lifts to see any significant benefit to service standards.

If customers use some personal queue knowledge, the service-standard improvements are smaller with each additional lift replacement. For example, as shown in Exhibit 4, with no new terrain, service standards start at .77 for the existing configuration under the current demand level and increase to .86 with the addition of the first new lift, .89 with the second added lift, and .90 with the third new lift. Additionally, the service standard increases from .77 to .86 if the terrain is expanded with no lift upgrades. But only marginal service-standard differences are seen between the expanded terrain and no expanded terrain service standards if lifts are upgraded.

Thus, if management assumes that customers use personal queue information, the single biggest improvement in service standard will come from either one lift replacement or expansion of terrain.

The use of resort queue information shows a similar pattern: service standards dramatically increase with the first lift replacement, and increase less with each subsequent replacement. The resort information provides the biggest service-standard increase (8 percentage points) for the first lift replacement compared to other scenario changes.

> Enhancing the resort's queue information provides the single biggest improvement in service standards, regardless of the environmental variables.

Inter-day demand smoothing. Regardless of the capacity changes made to the resort, the use of inter-day demand smoothing provides only marginal improvements to service standards. The data in Exhibits 4 and 5 indicate that the change in service standard ranges from none at all to a three-percentage point improvement when inter-day smoothing is used. A similar level of improvement is seen with queue information; slightly larger service-standard improvements are seen when queue information is used in conjunction with inter-day demand smoothing. For example, in Exhibit 4, with the existing configuration, the service standard improves by 13 percentage points with no smoothing and 15 percentage points with smoothing.

Interaction of variables. As demand grows from the existing levels, the service standard drops from the existing point, .74, to .71 for 5-percent demand growth and to .60, with 20-percent demand growth. If customers use no prior queue information, the resort would have to add three new lifts to maintain existing service standards with 5-percent or 20-percent demand growth. On the other hand, assuming customers use personal queue knowledge, the resort would need to add one new lift or expand terrain to maintain service standards near existing service standards of .77 for both growth scenarios. Knowledge of the lift status turns out to be important. If the resort displayed queue information, no other changes would be required to meet 5-percent growth, while either expanded terrain or a new lift would be needed to accommodate 20-percent growth. In this case, the service standard would be 10-percent above the existing service standard. Thus, the sign system gives improved service standards for the growth scenarios.

Exhibit 5 illustrates degradation in service standards with the desired customer-class mix and the three alternative growth levels. Regardless of the growth level and assumptions about prior queue knowledge, if the current customer-class mix changes to the desired mix (i.e., 15-percent more beginners and intermediate skiers), there are no possible capacity changes that will maintain the current service standard of .74 to .77. Only by using lift-queue information can the resort maintain the current service standard (.87) with management's desired customer-class mix. The resort must add two new lifts in the no-growth case and three new lifts and inter-day smoothing in the 5-percent growth case. With 20-percent growth, the existing service standard can not be maintained with the strategies considered here. In this case, the highest achievable service standard (.80) requires three new lifts and inter-day demand smoothing.

Enhancing the resort's queue information provides the single biggest improvement in service standards from the existing configuration, regardless of the environmental variables. Using queue information, improvements range from a minimum of 10 percentage points (desired customer-class mix and 20-percent growth) to 39 percentage points (desired customer-class and 5-percent growth).

Optimal service-level strategies. To achieve the highest service standard in any growth scenario and customer-class mix, the resort would have to implement the highest level of each strategy: replace three existing lifts, expand the terrain, install lift-queue information, and use inter-day demand smoothing. It is also apparent from the results that several other options greatly improved service standards. For example, replacing two lifts and using enhanced queue information improves the service standards to within 1 percent to 4 percent of the highest service-standard strategy, depending on the growth and class-mix scenario.

Managerial Implications

Should management execute all possible strategies to achieve the highest service standard? The answer depends on the costs associated with each strategy level and the service standards that customers desire. Each of the variables under management's control involves a considerable capital investment: $1,500,000 for an upgraded lift, $2,000,000 for expanded terrain (not including extensive environmental-impact reviews), and $500,000 for queue-information signs.

If managers assume that service-standard improvements lead to growth, and that their marketing efforts eventually change the customer-class mix, the results of the simulation dictate the following sequence of improvements. First, management should install or improve queue-information signs. The signs are the least-expensive investment and offer the largest single improvement in service standards. For the signs to function effectively, the queue data must be updated constantly so that the information is actually helpful to the customers. If, as management estimated, 50 percent of the customers use the queue information to choose their next lift, the resort's lift facilities will be used more efficiently if the signs are current. If more than 50 percent use the queue information, service standards will further improve without additional capital investment.[9] Furthermore, by monitoring queue lengths, managers have the ability to track service standards on a continual basis.

Second, managers should upgrade at least one lift and preferably two lifts. After the two upgrades, they should monitor their customer-class mix and overall demand to see whether shifts in use are actually occurring. A third lift should be upgraded after a 5-percent increase in demand growth.

Further service improvements have marginal benefits compared to the dollars invested. Managers should assess the trade-off between the cost of inter-day demand smoothing and expanding into new terrain. The simulation evaluated the effects of shifting 10 percent of the weekend skiers to the weekdays. This strategy has a cost, and it is unclear what type of price reduction is needed to move different proportions of customers to off-peak days. Obviously, there is a limit to how many people can be shifted to weekdays. Similarly, expanding into new terrain has a limited effect because all possible expansion terrain at this resort appeals to intermediate and advanced skiers. Consequently, if the desired customer-class mix shifts to less-skilled skiers (e.g., beginners), these skiers will not disperse to the new terrain. Because those two strategic objectives are in direct conflict with each other, the expansion strategy contributes little to the increase in service standards, while the customer-class mix that managers thought they wanted actually degrades service standards. Clearly, management should carefully evaluate the per-capita revenue gains from the desired customer mix versus the service-standard declines. For example, the family-skier segment increases the percentage of beginners and intermediates in the customer-class mix (which would degrade service) but, according to market-research data, this group spends more money on lessons and amenities than do experts. Balanced against that prospective increase is the potential erosion of long-term profits from skilled skiers who cannot tolerate the resulting degraded service standards. Managers will need to evaluate the potential effect of poor service standards versus the gains from increased revenues from the family (beginner) segment.

A caveat to our recommendations is our assumption that past performance is a good predictor of future information. Significant changes in the political, economic, or social environment would reduce the applicability of our simulation. For example, if changing weather patterns reduced the average snow pack at the resort, while increasing the snow pack at competitors' resorts, then our recommendations would not be valid.

The conflict of marketing and operational goals highlights the need for coordinated, holistic business strategies in hospitality firms. Our research suggests that a hospitality firm, faced with the objective of maximizing the number of customers entering a system subject to limits on waiting time, should prioritize the efficient use of existing facilities before looking at capacity expansion and external demand-management

[9] Obviously, if less than 50 percent of skiers use the sign information, lower benefits will accrue from the use of those signs. As noted earlier, however, the resort management's best *a priori* conservative guess was that half the skiers would use the information.

strategies. Another reason this advice is prudent is that one cannot truly judge the benefit of different capacity- or demand-management strategies if one is not currently optimizing the use of one's existing facility.

As an example of how these findings can be applied in other hospitality-service networks, consider the case of theme parks. Theme parks are similar to ski resorts: customers queue for various activities, including rides and food service. Long lines not only reduce customer satisfaction, but customers waiting in lines are not walking around spending extra, discretionary dollars. Thus, before adding new rides to lower wait times for existing rides, managers should look for ways to improve their existing line management. As in the case of our ski resort, theme park managers can provide signs that display up-to-the-minute wait times at various attractions, or institute a reservation system, where one does not need to join an attraction's line until one's reserved time. ■

Madeleine E. Pullman, Ph.D. (photo on left), is an assistant professor in the College of Business, Colorado State University (mellie.pullman@mail.biz.colostate.edu).
Gary M. Thompson, Ph.D. (photo on right), is an associate professor in the School of Hotel Administration, Cornell University (gmt1@cornell.edu).

© 2002, Cornell University