Labor Scheduling, Part 4

Controlling Workforce Schedules in Real Time

by Gary M. Thompson

Once you've built your employee work schedule, here's how can you ensure that it really works.

Deploying labor effectively is one of the most important tasks that front-line managers perform in service organizations. When a restaurant has too few employees, or its employees have the wrong skills, the result can be long lines in the foyer, long waits at the table, overworked employees, and poor service to customers. Having too many employees reduces operating margins, if extra hours are scheduled, or results in employees getting fewer work hours than they desire, if extra hours are not scheduled. The balancing act of best meeting customer demand while best satisfying employee requests—the process of deploying labor—is commonly known as labor or workforce scheduling. Hospitality services are typical in that labor often represents a large portion of the costs under managerial control. Controlling this cost, through effective scheduling, poses a challenging task to hospitality managers, particularly in light of employees' skills and requests, governmental regulations, company policies, and contractual obligations.

I have characterized workforce scheduling as comprising four tasks.[1] As explained in the previous articles

[1] G. Thompson, "Labor Scheduling Using NPV Estimates of the Marginal Benefit of Additional Labor Capacity," *Journal of Operations Management*, Vol. 13 (1995), pp. 67–86; and G. Thompson, "Assigning Telephone Operators to Shifts at New Brunswick Telephone Company," *Interfaces*, Vol. 13, No. 4 (July–August 1997), pp. 1–11.

***Gary M. Thompson**, Ph.D., an associate professor of operations management at the Cornell University School of Hotel Administration «gmt1@cornell.edu», has written a four-part series on labor management, of which this is Part 4. Parts 1 and 2 appeared in the October and December 1998 issues of* Cornell Quarterly, *and Part 3 was published in the February 1999 issue.*

in this series, the first task is to forecast customer demand for the service, the second is to translate the forecasts of customer demand into employee requirements, and the third is to develop a workforce schedule, using the employee requirements as inputs.[2] When those three steps are completed, a manager would have a forecast of the elements of the service transaction (particularly, customer arrival rates), a list of the number and skills of employees needed, and a specification of who is working where and when. The first three tasks are all planning activities in that they are conducted in advance of the service transactions. In contrast to the first three tasks, the final task involves the real-time control of the schedule, in which the manager assesses whether the schedule is ensuring that customers are actually being served as planned.

This fourth step, which involves comparing operating reality to the planned schedule, is the essential final piece to ensuring that your customers will be served appropriately. The difficulty in making sure service is as it should be occurs when real-time imbalances between labor capacity and customer demand arise. Such imbalances occur because demand rarely materializes the way one forecast it and because employees do not always perform the way one anticipates (e.g., they may be sick or late). The uncertainty in demand forecasts and employee performance highlights the need for effective real-time control that ensures that the *actual* schedule is effective.

In this paper I explain how a manager can assess with reasonable certainty whether the forecasted schedule is, in fact, matching the day's customer demand. I'll also touch on some of the actions a manager can take when demand does not match the forecasted schedule. In particular, I'll explain the value of having available cross-trained employees, particularly in those situations when demand is uncertain.

Because the approach I outline works best when customer counts are relatively high, this article will have the greatest applicability in high-volume restaurants and large hotels with substantial walk-in demand, as well as such other high-volume operations as reservation centers. The relatively high variability that occurs with low customer counts reduces one's ability to predict early in the day the likely business volume for that day. Most of my analysis is aimed at high-volume operations. However, hospitality services with low customer counts can use some of the techniques I mention, in particular those that are short-lived, as I explain next.

Real-time-control Actions

Real-time-control actions can be categorized either as short-lived or long-lived. Short-lived actions are those that affect only a small period of the operating day, typically a few minutes to an hour. Such actions are easily revocable. Short-lived actions include sending employees to or recalling them from break, extending the length of an employee's shift (including overtime), and asking employees to perform different tasks for a little while. Long-lived actions are those that will affect a period longer than an hour and entail a greater commitment of resources. They include sending employees home early, calling additional employees in to work, and reassigning employees to different jobs.

The key issue of real-time control is determining when to take an action that modifies the original schedule and (if the determination is positive) whether to take a short-lived or long-lived action. Short-lived actions have a relatively small effect on costs and on customer service, while long-lived actions can not only affect operations, but they can be difficult to reverse. Thus, for a manager to confidently take a long-lived action, she must be able to predict the hospitality operation's demand for that day. Say, for instance, that demand on Mondays is fairly consistent. With that consistency, a manager can make a statement like: "If we're slower than we anticipated by 11 o'clock, then it's likely we'll be slow for the whole day." The final step in scheduling is to be able to quantify that statement and to identify as early as possible whether a given day as a whole will be slower or busier than was forecast. A manager who can make that judgment can confidently take long-lived actions, such as sending workers home. Without that predictive confidence, however, the manager will have at her discretion only short-lived actions (e.g., assigning more employees to side work).

A Step-wise Approach to Tracking Demand

The bulk of this article gives a five-step process for predicting a day's customer counts. The steps are as follows: determine whether the operation enjoys consistent demand; identify the proportion of sales accruing to each planning period; categorize each day by its business volume; run a simulation of each day's business pattern to develop business-volume-consistency charts; and track customer counts against the simulation to predict day-end business volume. I'll explain each of these steps, although some steps will be familiar from the initial process of developing the demand forecast.

Step 1—**Determine the extent to which each day has a consistent demand pattern.** I addressed the issue of consistency in within-day demand in the first paper of this series. Consistent within-day demand means that each period within the day has a consistent portion of the daily demand, regardless of the total volume of business on a day. Exhibit 1 (which is a copy of an exhibit from that earlier paper), provides an example of this consistency.[3] In this step, the manager plots the sales in each planning period as a percentage of sales for the day. The graph shows that demand is similar on the four consecutive Mondays, building to a secondary peak around period 15, experiencing a lull through period 25, and then building to the primary peak in period 40, followed by a tapering off throughout the remainder of the day. Running a correlation of the daily data, as shown in Exhibit 2 (also from the first paper),[4] yields strong correlations of 0.79 and higher, which indicates consistent within-day demand.

Step 2—**Identify the proportion of daily sales occurring in each planning period.** Once a manager has established whether within-day demand is consistent, she needs to identify the proportion of daily sales that occur in each planning period (any given division of the day, but often a 15-minute period). One begins by calculating an average proportion of sales in each period and applying the smoothing technique I described in the first paper to obtain the de-

Exhibit 1
Same-day-sales comparisons, week to week

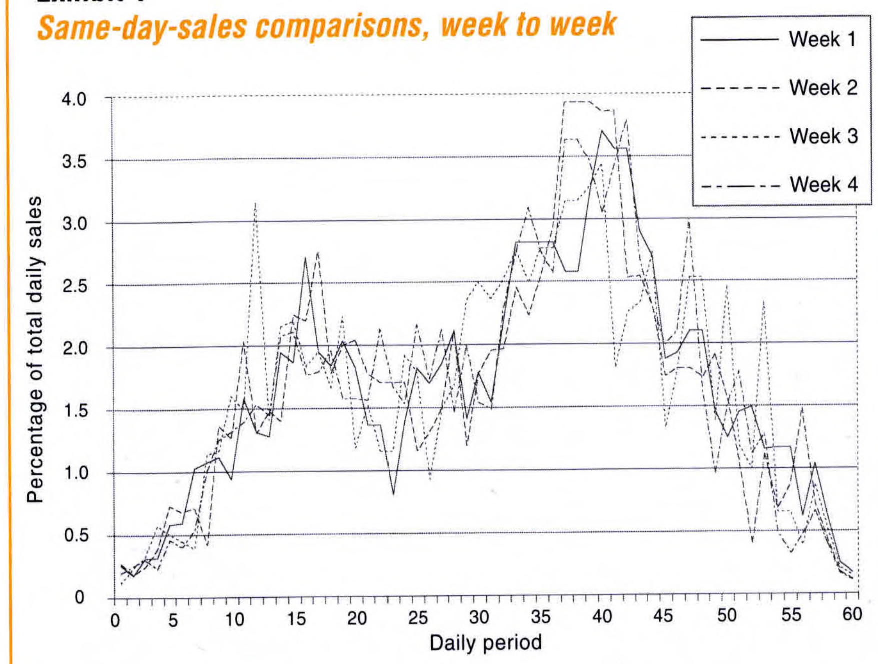

Sales are shown for every 15-minute period as a proportion of total daily sales, for a particular day of the week (e.g., Mondays), for four consecutive weeks.

Exhibit 2
Correlations between weeks for the data in Exhibit 6

	Week 1	*Week 2*	*Week 3*	*Week 4*
Week 1	1			
Week 2	0.8708	1		
Week 3	0.7954	0.7889	1	
Week 4	0.9039	0.8679	0.8262	1

The high correlation values, ranging from 0.79 to 0.90, indicate the applicability of an aggregation–disaggregation approach to forecasting within-day sales.

[2] G. Thompson, "Labor Scheduling, Part 1: Forecasting Demand," *Cornell Hotel and Restaurant Administration Quarterly*, Vol. 39, No. 5 (October 1998), pp. 22–31; G. Thompson, "Labor Scheduling, Part 2: Knowing How Many On-duty Employees to Schedule," *Cornell Hotel and Restaurant Administration Quarterly*, Vol. 39, No. 6 (December 1998), pp. 26–37; and G. Thompson, "Labor Scheduling, Part 3: Developing a Workforce Schedule," *Cornell Hotel and Restaurant Administration Quarterly*, Vol. 40, No. 1 (February 1999), pp. 86–96.

[3] See: Thompson, Part 1, p. 28.
[4] *Ibid.*

Exhibit 3
Comparing unsmoothed and smoothed demand forecasts

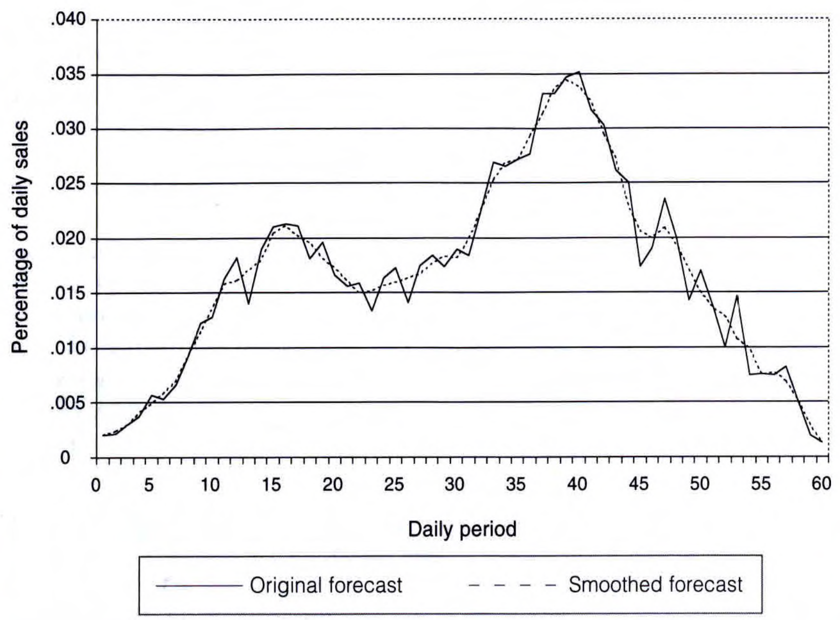

mand curve illustrated in Exhibit 3.[5] As discussed in the first paper, managers should be able to articulate the reasons customer demand materializes at the times it does. For example, given the nature of the service, there are reasons why the peaks and valleys in demand fall at the times they do.

Step 3—**Categorize each day according to its business volume.** The next step is to label a given day according to a customer-volume category. A reasonable way to do this is to establish, say, five categories of business volumes that cover the range of the operation's total daily customer counts. Level 1 would be the lowest level of business, while level 5 would be the highest.

For example, consider a hospitality service that typically serves 1,000 patrons on an average day, but where customer counts can range from 500 to 1,500. One could set up demand categories by dividing the thousand-customer spread into five even levels. Under that scheme, the five customer-volume levels would be 500 to 700 (level 1), 701 to 900, 901 to 1,100, 1,101 to 1,300, and 1,301 to 1,500 (level 5).

With those categories in mind, a manager would tally the actual customer demand by planning period over the course of the day. Such a graph is termed a realization. Exhibit 4 shows a hypothetical period-by-period realization of a level-1 day, during which a total of 598 customers are served. Note that the demand realization is generally consistent with the average business by period shown in Exhibit 3. There's a peak around period 15, a lull until approximately period 25, and then the facility hits its greatest demand around period 40. However, a key feature of Exhibit 4 is its variability, as represented by the sharp peaks and valleys from period to period. This variability—the randomness of customer demand—is the characteristic of service systems that is the prime driver of the need to make real-time capacity adjustments.

A more useful tool than the period-by-period realization of Exhibit 4, however, is a cumulative-business-volume realization. Exhibit 5 shows the cumulative realization for the customer-arrival data in Exhibit 4. Exhibit 5 shows that 100 customers were served by period 15; 200 customers were served by period 26; and 300 customers were served by period 34. A manager needs cumulative realization for the next step of the process, which is to assess the extent to which the realized demand is consistent with the forecasted demand.

[5] See: Thompson, Part 1, p. 29. In essence, the smoothing technique is a repetitive process of taking a mean of a given period and the two adjacent periods, and moving to the next period to take another three-period mean. This process shears off small blips in the demand graph.

Step 4—**Simulate realizations of the business-volume categories and develop business-volume-consistency charts.** To assess the fit between actual demand and forecasted demand, one needs a set of cumulative realizations against which to compare a given day's demand pattern. To get that set, the manager must either collect real data or simulate a set of cumulative realizations of these business volumes. One should collect or simulate over 100 realizations of each business volume (that is, 100 days of demand figures). Once these realizations are collected or simulated, one can develop business-volume-consistency charts that show the range of cumulative customer counts by period within a day. Because of the difficulty of collecting enough real data to develop the business-volume-consistency charts, I recommend that you develop a simulation.

Simulation is a useful tool for generating more data about operations than is readily available. Customer arrivals in the hospitality business typically follow Poisson distributions (rather than bell curves; see Exhibit 6 on the next page). With a Poisson distribution, one can develop an equation that specifies the typical arrival patterns. That equation is derived as follows: if a customer arrives at time t, then the next customer would arrive at time $t + [(-1 \div m \times \ln(R)]$, where $\ln(R)$ is the natural logarithm of a random number between zero and one, and m is the mean rate of customer arrivals. (One is calculating the negative reciprocal of m and multiplying it by the natural log of R. That mean rate of customer arrivals (m) will vary over the day based on the historical level of business. Using a series of random numbers, then, one can simulate customers' arrival times. Doing this one time is not especially useful,

Exhibit 4
Period-by-period realization of a day's customer demand (hypothetical)

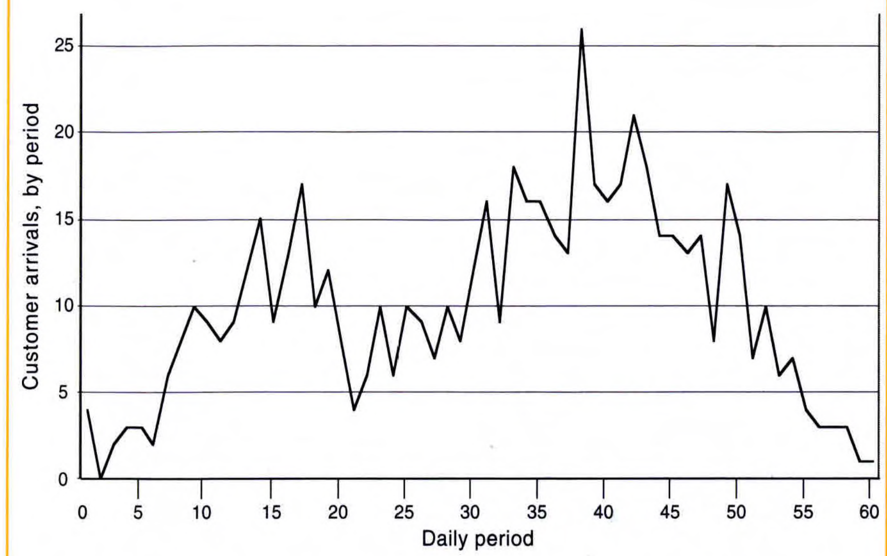

This graph represents a hypothetical period-by-period realization of a level-1 day. A key feature of the measurement is the variability, as represented by the sharp peaks and valleys from period to period, which represent the randomness of customer demand.

Exhibit 5
Cumulative realization of a day's hypothetical customer demand

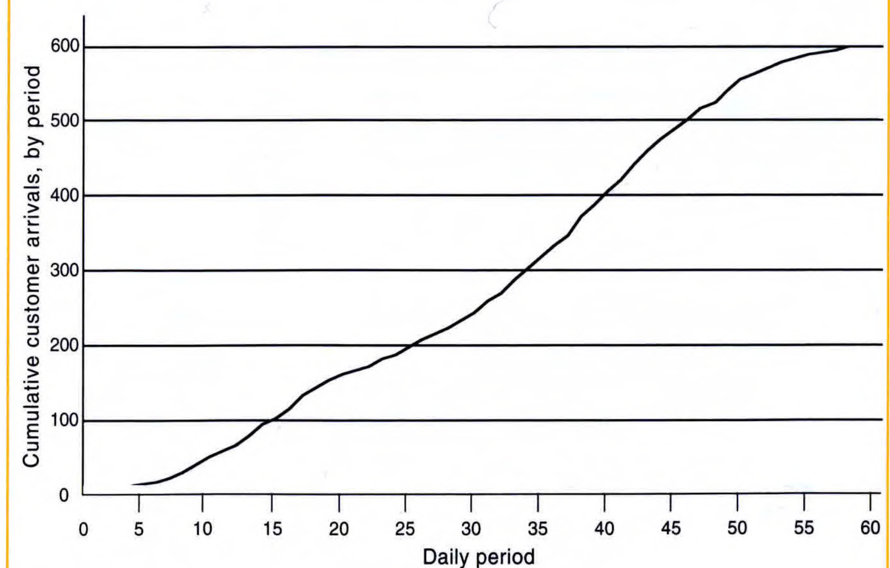

This graphic shows the cumulative realization for the customer-arrival data in Exhibit 4. A manager needs cumulative realization to assess the extent to which the realized demand is consistent with the forecasted demand.

Exhibit 6
Poisson distributions of customer arrivals

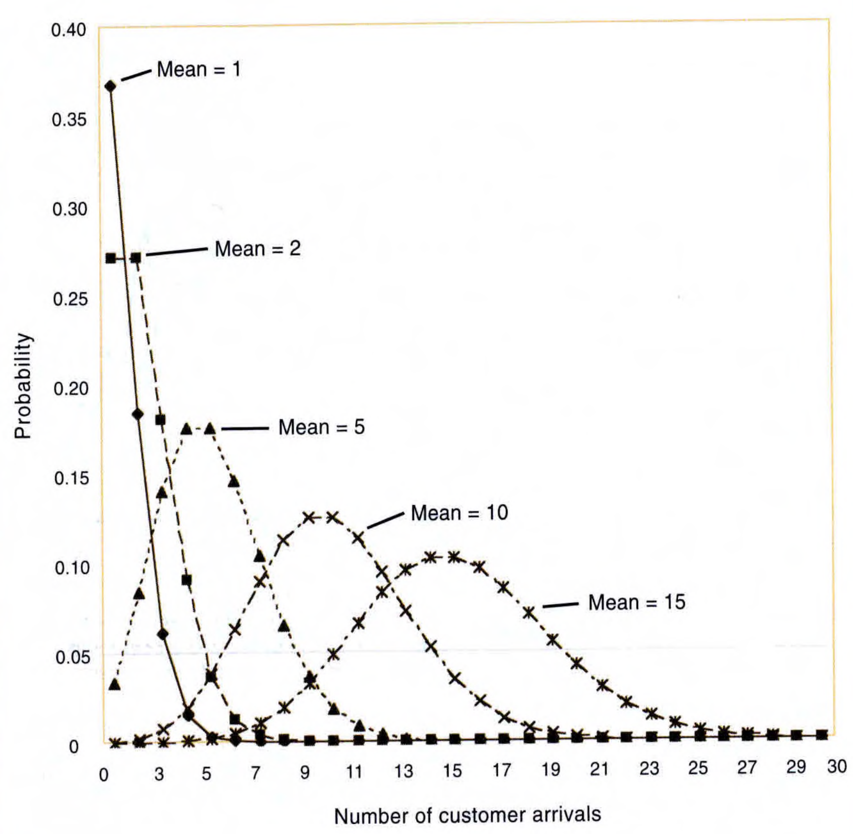

Customer arrivals in the hospitality business typically follow Poisson distributions (rather than bell curves). As such, one can develop an equation that specifies the typical arrival patterns (as described in the text). Running the calculation through many iterations gives one a reasonable simulation of arrivals for different business volumes.

but running the calculation through many iterations gives managers a reasonable simulation of customer arrivals for different business volumes.

Examples of business-volume-consistency charts are shown in Exhibits 7 and 8. Here's how they help you determine whether your schedule forecast is holding. Exhibit 7 shows level-1 volume (500 to 700 customers served for the day) based on 200 simulated realizations. The 100-percent line identifies the greatest number of customers served at any given point for days of a particular demand level (while the zero line indicates the fewest customers served at any point). The 50-percent line is the median number of customers served by that point, for a given day's demand. The 25- and 75-percent lines are respectively the first and third quartiles of the customer counts. Looking at period 30, and the 100-percent line, one will note that on any day that this operation served more than 307 customers by period 30, the operation *never* recorded a level-1 day. In other words, demand at that level by that time foretells more customers coming so that the day's demand will exceed 700. From the zero line we observe, by the same token, that on no day did we serve fewer than 188 customers by period 30 and still end the day in the level-1 business volume. Finally, the 50-percent line shows that 50 percent of the time, the customer count was 250 or lower by period 30. (In other words, the 50-percent line is the median.)

The business-volume-consistency charts operate in similar fashion for other levels of business. Exhibit 8, for example, is a business-volume-consistency chart for a level-2 day (total demand of 701–900). Again keying on period 30, observe that the operation never had served fewer than 272 or more than 405 customers at that point on a level-2 day, while 50 percent of the time it served 334 or fewer customers by period 30.

Although it is not readily apparent from Exhibits 7 and 8, the business-volume-consistency charts are not always distinct for separate volume levels. For example, consider a situation in which the operation had served 290 customers by period 30. This customer volume is within the range the operation experienced for level-1 days *and* within the observed range for level-2 days. A manager must be able to discern those differences, and that is the last step in the process, to wit, predict the day-end customer count.

Step 5—**Track customer counts and predict day-end business volume.** Given the overlap between the business-volume-realization curves, one might question how a manager could hope to distinguish, say, a level-1 day from level-2 day. Here's one way. The idea is to match a given day's customer counts with the appropriate consistency curve. One does this by recording a cumulative customer count early in the day and checking how that graph predicts the day will end up based on that count. As one tracks the cumulative customer count for a given day, one can compare the actual cumulative demand to the simulated-realization curves. Make a count of the total number of comparable cases in the simulated realizations and record the frequency with which each business volume contributed to the total comparable realizations. A comparable case is one where the simulated cumulative customer count in the previous period was equal to or less than the current customer count while at the same time the simulated customer count in the current period equals or exceeds the current customer count. For example, let's say we had served 14 customers by the end of the fifth

Exhibit 7
Business-volume-consistency chart (low-demand day)

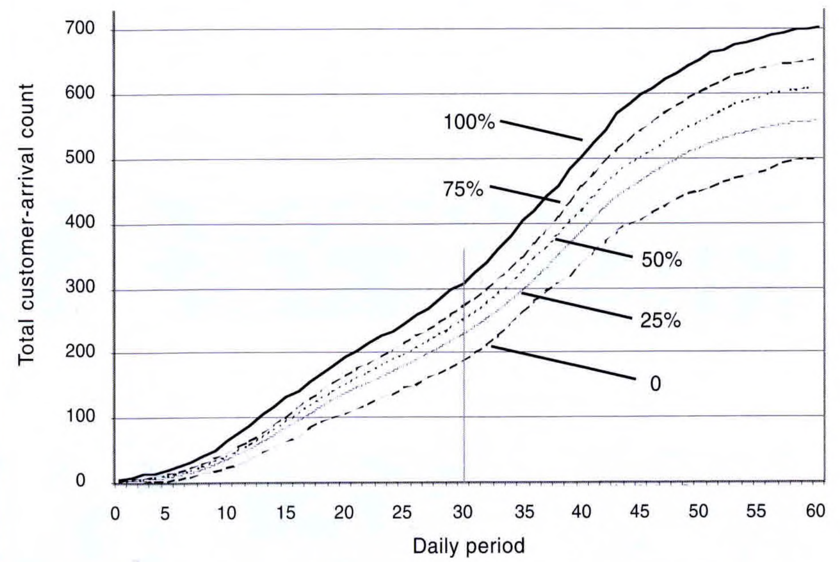

As explained in the text, Exhibit 7 shows level-1 volume (500 to 700 customers served for the day) based on 200 simulated realizations. The 100-percent line (top) identifies the greatest number of customers served at any given point for days of a particular demand level while the zero line (bottom) indicates the fewest customers served at any point.

Exhibit 8
Business-volume-consistency chart (medium-demand day)

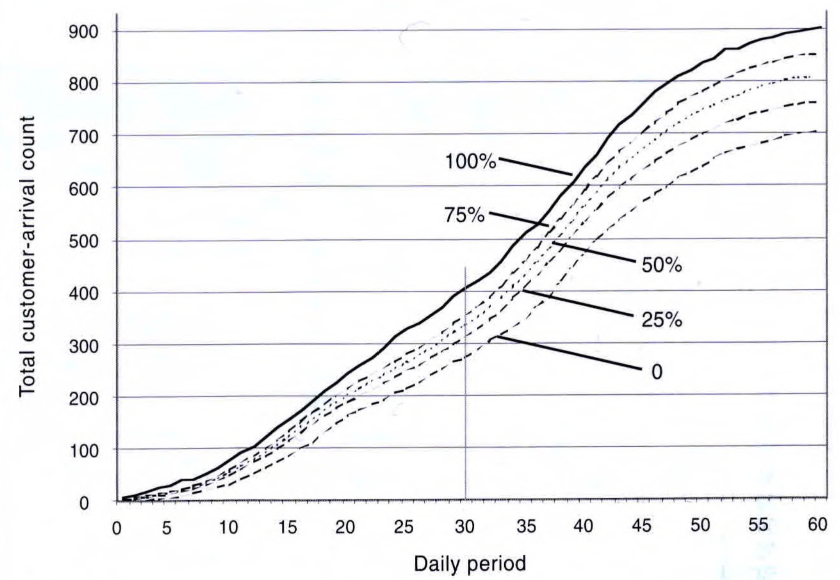

This chart represents a level-2 day (total demand of 701 to 900).

Exhibit 9

Sample use of customer counts to predict daily volume

Period	Cumulative Actual Customer Count	Number of Comparable Cases (percentage of total comparable cases) Level 1	Level 2	Level 3	Level 4	Level 5	Total Comparable Realizations
1	2	68 (25.56%)	69 (25.94%)	51 (19.17%)	43 (16.17%)	35 (13.16%)	266
2	7	185 (23.99%)	179 (23.22%)	150 (19.46%)	145 (18.81%)	112 (14.53%)	771
3	10	71 (16.55%)	100 (23.31%)	94 (21.91%)	94 (21.91%)	70 (16.32%)	429
4	13	63 (18.16%)	96 (27.67%)	78 (22.48%)	74 (21.33%)	36 (10.37%)	347
5	14	39 (31.97%)	46 (37.7%)	21 (17.21%)	12 (9.84%)	4 (3.28%)	122
6	21	117 (35.67%)	119 (36.28%)	62 (18.9%)	26 (7.93%)	4 (1.22%)	328
7	25	63 (35.80%)	74 (42.05%)	33 (18.75%)	6 (3.41%)	0 (0)	176
8	28	52 (61.90%)	28 (33.33%)	4 (4.76%)	0 (0)	0 (0)	84
9	37	115 (69.28%)	48 (28.92%)	3 (1.81%)	0 (0)	0 (0)	166
10	45	99 (75.00%)	30 (22.73%)	3 (2.27%)	0 (0)	0 (0)	132
11	53	93 (83.78%)	18 (16.22%)	0 (0)	0 (0)	0 (0)	111
12	65	119 (87.50%)	17 (12.5%)	0 (0)	0 (0)	0 (0)	136
13	73	74 (92.50%)	6 (7.5%)	0 (0)	0 (0)	0 (0)	80
14	87	102 (92.73%)	8 (7.27%)	0 (0)	0 (0)	0 (0)	110
15	103	101 (89.38%)	12 (10.62%)	0 (0)	0 (0)	0 (0)	113
16	116	73 (91.25%)	7 (8.75%)	0 (0)	0 (0)	0 (0)	80
17	128	62 (95.38%)	3 (4.62%)	0 (0)	0 (0)	0 (0)	65
18	147	82 (90.11%)	9 (9.89%)	0 (0)	0 (0)	0 (0)	91
19	155	33 (89.19%)	4 (10.81%)	0 (0)	0 (0)	0 (0)	37
20	171	54 (85.71%)	9 (14.29%)	0 (0)	0 (0)	0 (0)	63
21	186	38 (74.51%)	13 (25.49%)	0 (0)	0 (0)	0 (0)	51
22	198	32 (68.09%)	15 (31.91%)	0 (0)	0 (0)	0 (0)	47
23	204	23 (85.19%)	4 (14.81%)	0 (0)	0 (0)	0 (0)	27
24	215	33 (86.84%)	5 (13.16%)	0 (0)	0 (0)	0 (0)	38
25	221	14 (100)	0 (0)	0 (0)	0 (0)	0 (0)	14
26	229	22 (95.65%)	1 (4.35%)	0 (0)	0 (0)	0 (0)	23
27	239	33 (100)	0 (0)	0 (0)	0 (0)	0 (0)	33
28	248	27 (100)	0 (0)	0 (0)	0 (0)	0 (0)	27
29	259	30 (100)	0 (0)	0 (0)	0 (0)	0 (0)	30
30	270	27 (100)	0 (0)	0 (0)	0 (0)	0 (0)	27
31	288	42 (97.67%)	1 (2.33%)	0 (0)	0 (0)	0 (0)	43
32	307	41 (93.18%)	3 (6.82%)	0 (0)	0 (0)	0 (0)	44
33	332	48 (92.31%)	4 (7.69%)	0 (0)	0 (0)	0 (0)	52
34	341	21 (95.45%)	1 (4.55%)	0 (0)	0 (0)	0 (0)	22
35	352	23 (92%)	2 (8%)	0 (0)	0 (0)	0 (0)	25
36	368	35 (100)	0 (0)	0 (0)	0 (0)	0 (0)	35
37	387	40 (100)	0 (0)	0 (0)	0 (0)	0 (0)	40
38	412	50 (100)	0 (0)	0 (0)	0 (0)	0 (0)	50
39	443	53 (100)	0 (0)	0 (0)	0 (0)	0 (0)	53
40	464	38 (100)	0 (0)	0 (0)	0 (0)	0 (0)	38
41	485	31 (100)	0 (0)	0 (0)	0 (0)	0 (0)	31
42	498	21 (100)	0 (0)	0 (0)	0 (0)	0 (0)	21
43	513	24 (100)	0 (0)	0 (0)	0 (0)	0 (0)	24
44	531	27 (100)	0 (0)	0 (0)	0 (0)	0 (0)	27
45	543	15 (100)	0 (0)	0 (0)	0 (0)	0 (0)	15
46	550	8 (100)	0 (0)	0 (0)	0 (0)	0 (0)	8
47	557	10 (100)	0 (0)	0 (0)	0 (0)	0 (0)	10
48	571	20 (100)	0 (0)	0 (0)	0 (0)	0 (0)	20
49	576	7 (100)	0 (0)	0 (0)	0 (0)	0 (0)	7
50	582	6 (100)	0 (0)	0 (0)	0 (0)	0 (0)	6
51	594	12 (100)	0 (0)	0 (0)	0 (0)	0 (0)	12
52	604	11 (100)	0 (0)	0 (0)	0 (0)	0 (0)	11
53	607	3 (100)	0 (0)	0 (0)	0 (0)	0 (0)	3
54	615	7 (100)	0 (0)	0 (0)	0 (0)	0 (0)	7
55	620	5 (100)	0 (0)	0 (0)	0 (0)	0 (0)	5
56	624	3 (100)	0 (0)	0 (0)	0 (0)	0 (0)	3
57	628	3 (100)	0 (0)	0 (0)	0 (0)	0 (0)	3
58	633	3 (100)	0 (0)	0 (0)	0 (0)	0 (0)	3
59	636	5 (100)	0 (0)	0 (0)	0 (0)	0 (0)	5
60	636	3 (100)	0 (0)	0 (0)	0 (0)	0 (0)	3

Exhibit 10

Business-volume likelihood, example 1.

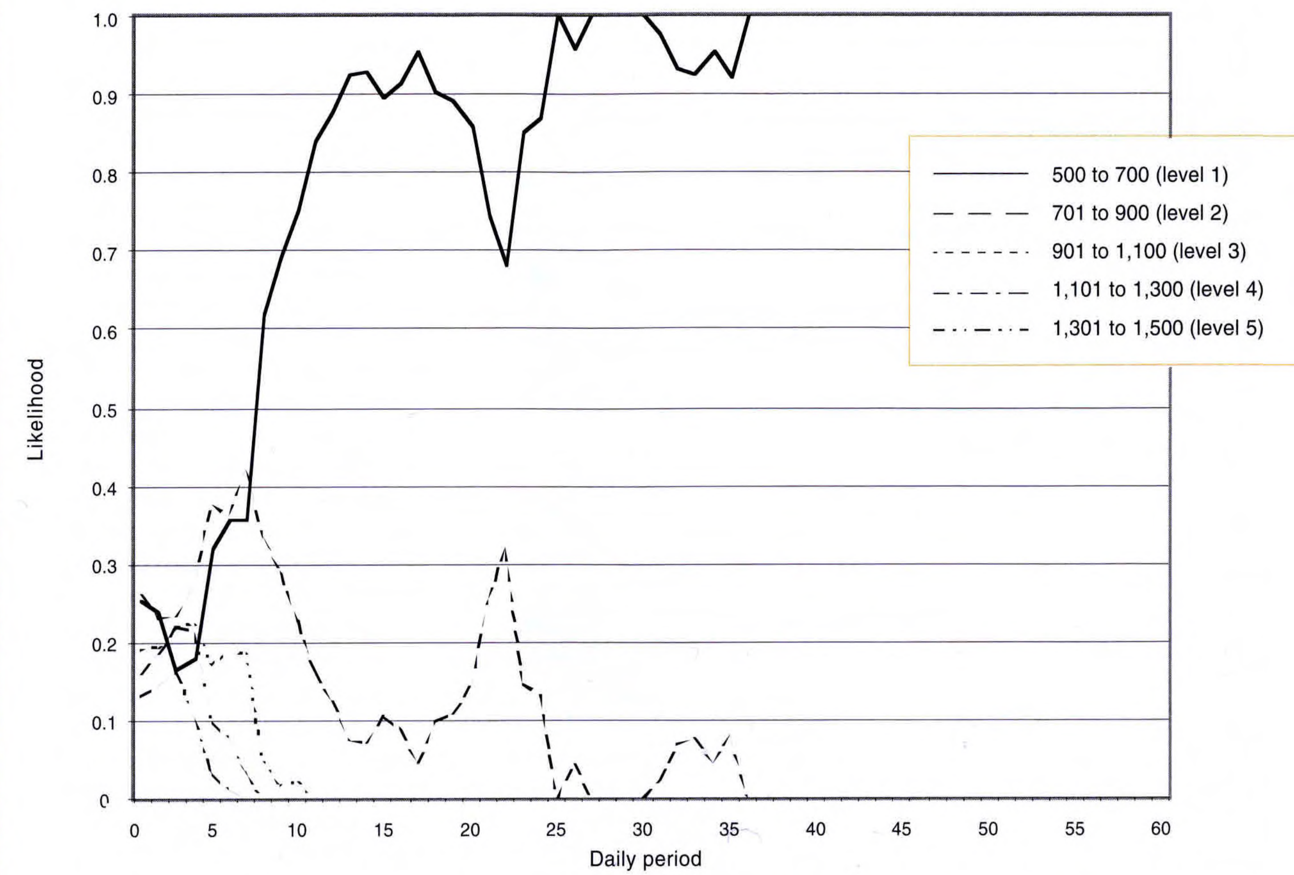

This graph displays the probabilities shown in Exhibit 9. The level-1 line, for instance, comes from plotting the 25.56 percent of period 1, the 23.99 percent of period 2, and so on.

period of the day. A comparable realization would be one in which 14 or fewer customers had been served by the end of period 4 and 14 or more customers had been served by the end of period 5.

As one moves through a day, one can monitor the level of business experienced to that point and compare it to the simulated realizations under different business volumes. Exhibit 9 shows an example of this approach. By the end of period 10, 45 customers had arrived on the particular day being tracked. Of the 200 realizations for level-1 business volume, 99 had customer counts that equaled or exceeded 45 customers *in period 10* and had 45 or fewer customers by the end of period 9. Level 2 had 30 comparative realizations, level 3 had three such realizations, and levels 4 and 5 saw no such comparable realizations. Thus, of the 132 comparable realizations, 75 percent resulted in a final daily customer count falling in level 1, just under 23 percent resulted in a final daily customer count falling in level 2, and slightly over 2 percent resulted in a final daily customer count falling in level 3. As of period 10, then, the manager has a strong indication that the final daily customer count will fall *below* level 3. By period 13 one can predict with a likelihood of over 90 percent that the day, as a whole, will experience level-1 demand.

The graph in Exhibit 10 displays the probabilities shown in Exhibit 9. The level-1 line in Exhibit 10, for instance, comes from plotting the 25.56 percent of period 1, the 23.99 percent of period 2, and so on. Showing the volume likelihoods graphically makes the day's demand easier to diagnose, or predict. Exhibits 11 and 12 illustrate other examples of how one might track such probabilities period-by-period throughout an operating day. Exhibit 11 shows that, as of period 20, there is greater than a 90-percent

Exhibit 11
Business-volume likelihood, example 2

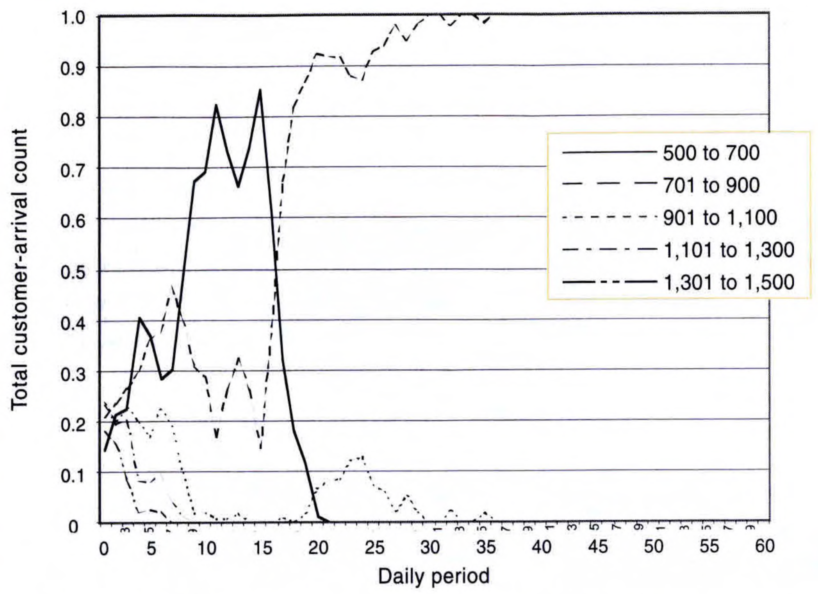

Showing the volume likelihoods graphically makes the day's demand easier to predict. There are a number of ways to track demand probabilities period-by-period throughout an operating day. In this example, as of period 20, there is greater than a 90-percent likelihood that the day, as a whole, will fall in business-volume-level 2 (701 to 900).

Exhibit 12
Business-volume likelihood, example 3

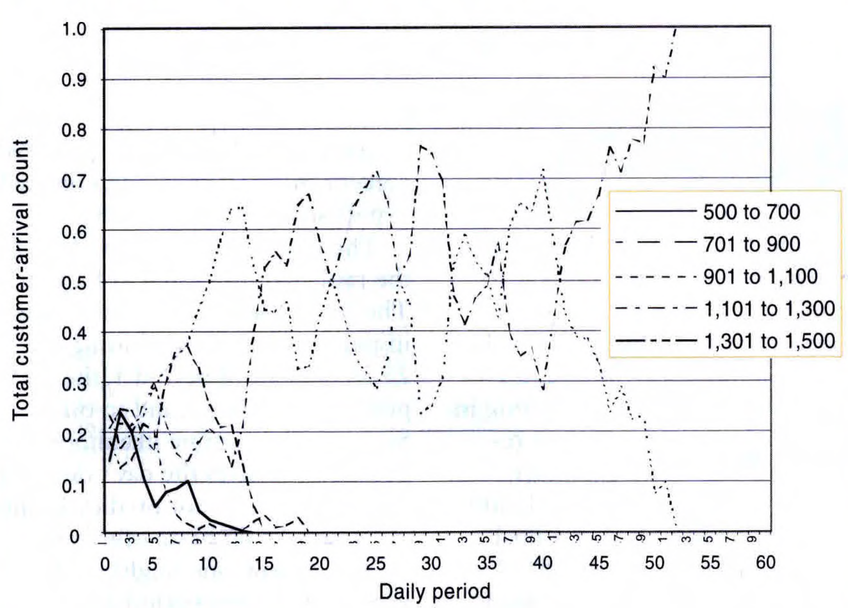

The customer-demand data illustrated here give a strong early indication (by period 15) that the day will hit level 3 or level 4 (901 to 1,300). However, it is not until period 50 that the indication becomes clear that the day's demand will end up in level 4 (1,100 to 1,300).

likelihood that the day, as a whole, will fall in business-volume level 2. The customer-demand data illustrated in Exhibit 12 give a strong early indication (by period 15) that the day will hit level 3 or level 4. However, it is not until period 50 that the indication becomes clear that the day's demand will end up in level 4.

Exhibits 10, 11, and 12 raise the question of why one can predict the day-end business volumes earlier on some days than on others. The answer lies in the fact that initial customer counts may fall near the break points between categories. Turned around, the point is that the closer the final customer count is to the breakpoints between categories, the longer it takes to predict the final daily demand. Thus, if the customer counts are right on the cusp of two categories, the manager might not be able to establish the final daily demand until the end of the day. However, choosing one or the other of two adjacent categories is not the point of this process, so much as getting an early indication that the customer count will fall in one or the other of two adjacent volume categories. As in the case of Exhibit 12, one would have a consistent early signal that the day will likely be in one or the other of two adjacent volumes. In this case, the manager still can make the necessary real-time schedule adjustments, even if the certainty of the outcome isn't great.

The converse is also true: the stronger the indications are that final demand will fall in the middle of a category, the earlier in the day one can predict that day's business volume. Similarly, extreme volumes (i.e., level 1 or level 5) will be easier to predict than mid-range volumes.

Real-time-control Actions

To develop the historical baseline data needed for this procedure, a manager should periodically perform steps 1 through 4. In contrast, step 5 should be performed hourly or even more frequently, because it is the monitoring step that allows one to predict day-end business volumes. In turn, predicting total daily business early in a day allows a manager to take appropriate long-lived actions to adjust employee schedules. In the second article in this series, I explained how uncertain demand (and an easier ability to send employees home than to call them in) causes managers to increase staffing levels.[6] Thus, even with an expectation of average (level-3) demand, a manager probably would develop a schedule based on a level-4 volume just to be sure that all customers will receive appropriate service. Then, if the manager gets a strong early indication that demand will fall into a lower volume category, she can take appropriate long-lived actions (e.g., asking for volunteers to go home without pay).

If, by contrast, she has set a level-4 schedule and then gets a strong signal that demand will hit level 5, she would want to take long-lived actions like extending employees' shifts, offering overtime, and perhaps calling additional employees in to work. She might even consider reoptimizing the day's labor schedule based on the new demand information.

On the other hand, if the operation is experiencing a real-time capacity-demand imbalance, but the indication is not clear as to what volume the day will see, the only valid real-time-control actions are short-lived (e.g., reassigning employees). Taking long-lived actions runs the risk of the need for further (and unnecessary) actions later in the day.

The Value of Cross-training

Even a relatively solid forecast contains the possibility of error, which is why managers usually err on the safe side and overstaff their operations. One way of reducing the effect of forecast uncertainty is by employing cross-trained workers. By having a cadre of cross-trained employees a manager gains scheduling flexibility, because she can deploy her cross-trained workers where they are most needed. Instead of counting the number of bussers and the number of runners, for instance, the manager could cross-train people for multiple jobs (including table servers) and set the schedule according to an estimate of the total help needed on the floor.

As an illustration of the value of cross-trained employees, consider the following example. Say that an operation has three different positions. Each of the positions would ideally be staffed by a complement of 10 employees, if the demand forecast were perfectly accurate. Exhibit 13 shows that with that perfect demand forecast the hourly cost of the system—both labor costs and the cost of customer waiting—would be \$335.71.[7] As the uncertainty in the demand forecast increases (as measured by the coefficient of variation of forecast

[6] Thompson, Part 1 (October 1998), p. 35.

[7] For a discussion of how to calculate the cost of customers' waiting time, see: Thompson, Part 2 (December 1998), p. 32.

How real-time control might work in a theme park

The principles explained in the accompanying article were developed in restaurants, hotels, and theme parks. As I stated in the main text, a real-time-control system (RTCS) is well-adapted to any service establishment that has high customer counts, including theme parks. In the case of a theme park, the RTCS would receive customer-arrival data from the park gates. Based on the day's weather, the RTCS would continually update its prediction about the business volume to be experienced throughout the day. The RTCS would also be fed real-time information from the payroll system—tracking which employees are late, or who have called in sick, for example. Finally the RTCS would receive real-time information from all point-of-sale systems within the park and from other data-tracking devices, such as queue-length monitors.

Using the current—and predicted—business volumes, the RTCS would serve as a management-decision aid: reoptimizing the labor schedule for the remainder of the day, recommending when to call extra employees in to work, when to send employees home, when and which employees to switch between positions to maximize the benefit to the organization, and when to send or recall employees from breaks. With complex hospitality service systems, like theme parks, RTCSs are the last, and presently uncharted, frontier of good labor management.—*G.M.T.*

Exhibit 13

A comparison of staffing allocations with and without cross-trained employees

COV*	*Base Case* (no cross training)		*Scenario 1* (20% wage & benefit premium for cross-trained employees)				*Scenario 2* (10% wage & benefit premium for cross-trained employees)			
	Staffing level**	Hourly cost	Staffing level**	Hourly cost	Hourly savings	Percentage savings	Staffing level**	Hourly cost	Hourly savings	Percentage savings
0.25	13/13/13/**0**	$403.81	9/9/9/**6**	$371.54	$32.27	7.99%	9/9/9/**6**	$365.54	$38.27	9.48%
0.20	12/12/12/**0**	378.51	10/10/10/**3**	360.41	18.10	4.78%	9/9/9/**5**	355.41	23.10	6.10%
0.15	11/11/11/**0**	357.59	10/10/10/**2**	349.77	7.82	2.19%	9/9/9/**4**	346.69	10.90	3.05%
0.10	10/10/10/**0**	347.62	10/10/10/**1**	341.26	6.36	1.83%	10/10/10/**1**	340.26	7.36	2.12%
0.05	10/10/10/**0**	335.71	10/10/10/**0**	335.71	0	0	10/10/10/**0**	335.71	0	0
0.00	10/10/10/**0**	333.12	10/10/10/**0**	333.12	0	0	10/10/10/**0**	333.12	0	0

* Coefficient of variation of the forecast error.

** Best staffing levels are indicated as follows:
number of people in position 1 / number of people in position 2 / number of people in position 3 / **number of cross-trained employees.**

error), the ideal staffing level increases to 13 employees per position and the total hourly cost rises to $403.81.

Exhibit 13 also shows the effect on hourly costs and staffing decisions when a manager can draw from a pool of cross-trained employees. Exhibit 13 considers scenarios where the cross-trained employees receive pay premiums of 20 percent and 10 percent compared to the standard employees. Assuming a 20-percent wage-and-benefit premium for the cross-trained employees, the ideal allocation of employees under the highest level of forecast inaccuracy would be to assign nine employees to each of the three positions and have six cross-trained employees who would be assigned in real time to the positions so as to balance the workload. This labor allocation would require 33 employees in total and cost $371.54 per hour, representing a 15-percent reduction in the number of employees and an 8-percent cost saving compared to staffing the positions with dedicated employees.

A close examination of the results in Exhibit 13 reveals several patterns. First, without cross-trained employees, higher forecast inaccuracy leads to **(a)** higher staffing levels and **(b)** higher hourly costs. Second, using cross-trained employees, higher forecast inaccuracy yields **(a)** a larger number of cross-trained employees and lower numbers of position-specific employees and **(b)** greater savings from cross-trained employees. Finally, when the cross-trained employees are relatively less expensive (than regular workers), **(a)** more cross-trained and fewer position-specific employees are warranted and **(b)** larger savings accrue from having cross-trained employees. Indeed, under the highest level of forecast inaccuracy, having a pool of cross-trained employees reduced employee needs by 15 percent and reduced costs by over 9 percent. Although it is not shown in Exhibit 13, the benefit of cross-trained employees is also greater when the employees can be shared across more than two jobs.

The availability of a cross-trained labor pool should be incorporated during the development of a labor schedule (Part 3 of this series). However, the actual deployment of the cross-trained labor would occur in real time, when employees would be assigned to the positions most useful to the hospitality firm. CQ