Planning-interval Duration in Labor-shift Scheduling

by GARY M. THOMPSON

This article proposes guidelines for managers and researchers regarding appropriate planning intervals in labor scheduling. This study distinguishes between planning intervals used for data collection and intervals used for scheduling employees. This distinction is relevant because it is not necessary for the two interval types to be of the same duration. The results of a simulation experiment show that both increased variability in the customer-arrival rate and high mean customer-arrival rates work to reduce the ideal length of both data-collection and schedule-development intervals. In fact, with a mean arrival rate of five customers per minute, the most profitable labor schedules originated with a five-minute data-collection interval combined with a five-minute schedule-development interval. At present, because of the paucity of procedures that can accommodate such brief intervals, managers will find fifteen-minute data-collection and schedule-development intervals generally to be the most effective.

Keywords: labor scheduling; planning

Labor scheduling—the process of matching the number of employees working with the number of employees needed to provide the desired level of service over the operating day—is frequently a large determinant of a hospitality organization's efficiency. Having too few employees on hand jeopardizes customer service and may result in the loss of present and future sales. Conversely, having too many employees on hand is inefficient, since the customer-contact activities performed by service personnel cannot be inventoried. Compounding the difficulty in developing labor schedules is the variation in customer demand that is common in restaurants.

In a series of articles published in this journal, I described the four main tasks that constitute workforce scheduling. Those tasks are forecasting customer demand,¹ translating the demand forecasts into staffing requirements,² developing the schedule,³ and controlling the schedule in real time.⁴ Labor scheduling usually requires that the operating day be divided into

small planning intervals, with the desired staffing levels determined for each interval. The question facing managers (and the one I investigate in this article) is, What is the best duration for the planning intervals? Planning intervals in the workforce-scheduling literature have ranged between fifteen and sixty minutes, but there is little prescriptive guidance for managers as to what length of planning interval is appropriate in particular environments.

Although customer arrival data may be collected by the hour, a fifteen-minute planning interval is still best for scheduling employees.

My interactions with hospitality managers and with vendors of labor-scheduling solutions indicate that operators most commonly use planning periods of fifteen minutes. A review of the academic literature, however, shows that thirty-minute and sixty-minute intervals are examined more commonly than are fifteen-minute intervals. This discrepancy concerns me since it may indicate that the academic literature is not addressing relevant problems. Certainly, using relatively long periods makes it easier to develop schedules since the alternatives that must be considered multiply when one uses short planning periods. However, if the hospitality industry is using fifteen-minute intervals, one could reasonably argue that academics should also work with intervals of that length. The discrepancy between the two groups' planning intervals motivated me to conduct the investigation reported in this article.

Break time. In a study with M. E. Pullman, I showed that reliefs—fifteen-minute

breaks—should be scheduled in advance, rather than only in real time (as has commonly been argued).⁵ This finding requires that schedules be developed using planning intervals of no more than fifteen minutes. However, there is no reason that data on customer arrivals have to be collected using the same intervals as those used in schedule development. For example, data on customer arrivals can be collected using sixty-minute planning intervals, while schedules are developed using fifteen-minute intervals. All one need assume is that each of the four fifteen-minute intervals in an hour has the same arrival rate (equal to one-quarter of the arrival rate for the hour). To my knowledge, this article is the first to make the distinction between these two types of intervals. Throughout the article, I will refer to these contrasting intervals as “data-collection intervals” and “schedule-development intervals.”

I seek to identify the best length for the data-collection and schedule-development intervals. In so doing, I attempt to answer the question of whether fifteen-minute (or shorter) planning periods are preferable, or whether it is appropriate to use thirty- and sixty-minute periods as has been done commonly by academic researchers. In the course of this discussion, I seek to identify the conditions under which different planning-interval lengths work best, so that I can provide prescriptive guidelines for managers and researchers choosing a planning-interval duration.

To distinguish among performance outcomes, I use the criterion of profitability attributable to effectively deployed labor. I examine the use of sixty-, thirty-, and fifteen-minute data-collection intervals, each with a fifteen-minute schedule-development interval and five-minute data-collection and schedule-development

intervals. In the next section, I describe the experiment I conducted to investigate planning-interval durations. I then discuss the results, including offering suggestions for future research.

The Experiment

This section describes the experiment I conducted to investigate the effectiveness of the various data-collection and schedule-development intervals. In the sections that follow, I address environmental experimental factors, managerially controllable experimental factors, problem commonalities, simulation details, and sources of information inaccuracy.

Environmental Factors

Environmental factors are characteristics of a hospitality environment that are outside a manager's control. In conducting an experiment, it is always prudent to select a wide variety of environmental factors since that helps ensure the generalizability of one's findings. The following five environmental factors figured into my experiment: the true arrival-rate pattern (four levels); the variability in the arrival-rate pattern (five levels); the length of operating day (two levels); the mean arrival rate (two levels); and the mean duration of customer service (two levels). Below I describe each of these factors.

Arrival rates. Customer-arrival-rate patterns simply chart the frequency with which customers arrive, by time of day. In selecting customer-arrival-rate patterns, I wanted to consider patterns that, via manipulation, would exhibit a wide range of variation. I selected four patterns, each with many peaks and valleys in demand. These patterns appear to be random, but more regular patterns (such as a trimodal

pattern that has peaks at meal periods) cannot be manipulated to exhibit the range of variability I wish to examine.

The magnitude of the peaks and valleys of customer arrivals—in other words, the change in the frequency of customer arrivals—was controlled by the second environmental factor, variability. I define a variability index as follows:

$$\text{Variability Index} = \frac{\sum_{n=2}^N |\lambda_n - \lambda_{n-1}|}{(n-1)\lambda}, \quad (1)$$

where

λ_n = true average customer arrival rate during the n^{th} five-minute interval; and

λ = average customer arrival rate across all intervals.

The variability index represents the average proportional change in mean arrival rates from one five-minute time block to the next across the operating day. The five arrival-rate levels had variability indices of 0.01, 0.03, 0.06, 0.09, and 0.12, corresponding to a range of low to high variability. Low variability can occur in hospitality organizations that use appointment systems to manage demand (restaurants, for instance). High variability can occur, for example, in telephone-call centers where demand is spurred by television advertisements or in quick-service restaurants (QSRs) that can be suddenly inundated by motor coach tourists who are taking a lunch break from traveling on a superhighway.

Open for business. The third environmental factor defines the length of the operating day, as tallied by the number of

hours that a hospitality firm is open for customers each day. For the two levels, I selected operating days of fourteen and twenty hours. As I show below, many more shift alternatives are available under the twenty-hour day than under the fourteen-hour day, making it more complicated to generate solutions to problems for the long day, as compared with the short day. Exhibit 1 illustrates the eight instantaneous customer-arrival-rate functions for a fourteen-hour operating day arising from the combinations of the four arrival-rate patterns and the lowest and highest levels of arrival-rate variability.

The fourth environmental factor is the mean customer-arrival rate. I selected mean arrival rates of two and five customers per minute. High mean customer-arrival rates, which are often seen in call centers and busy QSRs, generally mean increased staffing requirements.

Service. The final environmental factor, with two levels, is the mean duration of service. I selected durations of one minute (such as might be observed in a QSR) and ten minutes (such as might be observed in a call center for travel services). A service duration of ten minutes greatly increases the staffing requirements compared with the one-minute duration. Long service durations also increase the spillover effect of customer service, which occurs when service that commences in one planning interval spills over into subsequent periods. This spillover of customer service makes the task of determining staffing levels more problematic when one uses short planning intervals. I employed the procedure that I developed in 1993 to adjust for the spillover effect.⁶

The combinations of environmental factors in the full-factorial design resulted in a total of 160 test scenarios ($4 \times 5 \times 2 \times 2$).

For each of these scenarios, I simulated the collection of data and developed shift schedules using the four combinations of data-collection and schedule-development intervals, thus yielding a total of 640 observations in the experiment.

Managerially Controllable Factors

The managerially controllable factors—which are the decisions that I want to evaluate—concern the most appropriate length for data-collection and schedule-development intervals. I use four combinations of those intervals. Since I wish to schedule fifteen-minute reliefs (as I indicated above), the schedule-development intervals must be fifteen minutes or shorter.⁷ However, the data-collection intervals can exceed fifteen minutes. This may be preferable under circumstances where the customer-arrival rate changes little and when using short data-collection intervals could lead to inaccurate estimates of the number of customer arrivals (and, consequently, inappropriate staffing levels). That is why I evaluated the use of fifteen-minute schedule-development intervals with sixty-, thirty-, and fifteen-minute data-collection intervals.

Although fifteen minutes represents the shortest planning interval reported in the literature, I also wish to evaluate whether even shorter periods might be effective under certain conditions. Since I schedule fifteen-minute reliefs, the schedule-development intervals should be an even divisor of fifteen. I thus selected the combination of five-minute data-collection and schedule-development intervals as the fourth combination.

As a rule of thumb, the planning periods should be the shorter of the shortest break being scheduled (in this case, fifteen

Exhibit 1:
Eight Instantaneous Customer-arrival Rate Functions for Four Arrival Patterns

minutes) and the length of the periods necessary for adequate tracking and forecasting of customer demand. This means that one need not examine schedule-develop-

ment intervals shorter than fifteen minutes unless the data-collection intervals are also shorter than fifteen minutes. For example, if the break length is fifteen min-

Exhibit 2:**Number of Shift Alternatives under the Operating Day Lengths and SDI Lengths**

<i>Length of Schedule-development Intervals (minutes)</i>	<i>14</i>	<i>20</i>
15	67,371	134,211
5	10,543,731	21,404,859

utes long, there is no need to use schedule develop periods shorter than 15 minutes unless one is collecting data using those shorter intervals.

Problem Commonalities

Problem commonalities are the items I used across all experimental problems. Many of these items can be considered as managerially controllable, but they are simply outside the scope of the current experiment.

I used the labor-scheduling paradigm New Formulation of the Labor Scheduling Problem (NFLSP), which I presented in 1995.⁸ NFLSP is simply a way of representing the labor-scheduling problem in a profit-oriented way.

Restrictions. The operative restrictions on allowable shifts that I used were as follows: shifts were between six and nine hours in length, including an hour-long, unpaid meal break. The meal break was to be preceded and followed by at least 2.25 and no more than 5 hours of paid time. Two paid, fifteen-minute reliefs are required, one in the premeal work period, the other during the postmeal stretch. Each relief must be preceded and followed by at least 1 hour and no more than 3.75 hours of work. Exhibit 2 identifies the number of distinct shifts defined by these

timing restrictions. Particularly massive is the number of individual shifts under the longest operating day and five-minute planning periods.

I generated and solved the variants of NFLSP using commercially available software packages (GAMS⁹ for generating the problems and OSL¹⁰ for solving them) on the equivalent of a Pentium III 500-based personal computer. Solution times were limited to forty-five seconds, and I used the best integer solution found during that time if a verified optimal solution was not found. Integer solutions ensure that whole numbers of employees are scheduled for each shift and thus are realistic employee schedules.

In several instances of five-minute schedule-development intervals, the computer could not find any integer solutions within the forty-five-second time limit. For these problems, I doubled the available running time to ninety seconds. To represent shifts, I used a variant of the implicit model that I developed in 1995,¹¹ since finding solutions to an explicit model would be impractical, even for fifteen-minute schedule-development intervals. Unfortunately, the memory required for the scheduling problems with five-minute data-collection intervals under the twenty-hour operating day exceeded my system's capacity, and I could not solve those problems.

Information Inaccuracy

My experiment was designed to mimic the information inaccuracy typical of hospitality organizations. The rationale for this approach is that it makes little sense to evaluate planning-interval durations only under conditions of perfect information, because hospitality organizations do not operate under such conditions. Information inaccuracy arose in my experiment from the following sources. There was a limited amount of historical data with which to estimate the true customer-arrival rates; the variability in daily mean customer-arrival rates was unpredictable; estimates of the benefits and costs of good and poor service were inaccurate; and the relationship between customer satisfaction and their waiting times was hard to gauge.

Results

Exhibit 3 summarizes the number of problems for which each combination of data-collection and schedule-development intervals yielded the most profitable schedules. For the fourteen-hour day, the greatest number of most profitable solutions was yielded by the following combinations:

1. Sixty-minute data-collection intervals with fifteen-minute schedule-development intervals at variability-index levels of either 0.01 or 0.03;
2. Thirty-minute data-collection intervals with fifteen-minute schedule-development intervals at the variability-index level of 0.03 (a tie with number 1);
3. Fifteen-minute data-collection and fifteen-minute schedule-development intervals at variability-index levels of 0.06, 0.09, and 0.12; and
4. Five-minute data-collection and five-minute schedule-development intervals at a variability-index level of 0.12 level (a tie with number 3).

Simulation Details

The simulation experiment contains some simplifying assumptions, so that I could avoid modeling all the complications of a typical hospitality business. Those assumptions are as follows: (1) the schedule was implemented as developed (e.g., there were no real-time changes to the schedule); (2) the employees performed as scheduled (e.g., there was no absenteeism); (3) no customers balked from the queue; and (4) customer characteristics were stationary over the simulation period (e.g., there was no trend or cyclicity of the *mean* daily customer-arrival rate).

Data on customer-arrival times and the random numbers associated with customer-service times were generated for each of the eighty instantaneous customer-arrival-rate functions (calculated thus: 4 Arrival Patterns $\times$ 5 Levels of Arrival Rate Variability $\times$ 2 Operating Day Lengths $\times$ 2 Mean Arrival Rates). To control variance, the same data were used with both combinations of mean service time, yielding a total of 160 test problems. For each of these test problems, employee requirements were determined and optimal schedules were generated using NFLSP for the four combinations of time intervals, thus resulting in 640 data points.

For each test problem, sixty days of customer arrivals were simulated. Of these, ten days were deemed "historical" and used in setting NFLSP's profitability coefficients. These data were recorded in the form of the number of customers arriving in each historical data-collection interval (i.e., 1,680 five-minute, 560 fifteen-minute, 280 thirty-minute, and 140 sixty-minute intervals for the fourteen-hour day). In the situations where the data-collection intervals were longer than the schedule-development intervals, the mean customer-arrival rate was divided equally by the number of schedule-development intervals in each data-collection interval. For example, with a sixty-minute data-collection interval and a fifteen-minute schedule-development interval, a mean historical arrival rate of forty customers for an hour would be divided into arrival rates of ten customers in each of the four fifteen-minute schedule-development intervals in the hour. The profitability coefficients of NFLSP were determined following the process that I identified in a 1995 article,¹² using the arrival rates in each schedule-development interval.

The remaining fifty days of data were used to simulate the operation of the hospitality system with the generated schedule in effect. The average schedule profitability was recorded from the fifty days of simulated future operation of hospitality system. The fifty-day evaluation period was used only to yield a good estimate of the average profitability of implementing the schedule on a single day and does not represent the use of the same schedule for a fifty-day period. — G.M.T.

Exhibit 3:**Number of Most Profitable Schedules under Various Schedule Assumptions**

Day Length	Arrival-rate Variability Index	Data-collection Interval and Schedule-development Interval			
		60–15	30–15	15–15	5–5
14 hours	0.12	0	2	7	7
	0.09	0	1	10	5
	0.06	2	4	7	3
	0.03	5	5	4	2
	0.01	7	2	5	2
	Total	14	14	33	19
20 hours	0.12	0	3	13	
	0.09	0	3	13	
	0.06	0	6	10	NA ^a
	0.03	3	7	6	
	0.01	4	9	3	
	Total	7	28	45	

a. Too large to solve with my hardware and software combination.

For the twenty-hour day, the greatest number of most profitable solutions was yielded by:

1. Thirty-minute data-collection intervals with fifteen-minute schedule-development intervals at the variability-index levels of 0.01 and 0.03; and
2. Fifteen-minute data-collection and fifteen-minute schedule-development intervals at variability-index levels of 0.06, 0.09, and 0.12.

Exhibit 4 illustrates the mean profitability of the schedules by length of operating day and variability-index level. For the fourteen-hour day, five-minute data collection and schedule-development intervals yielded the most profitable schedules under the highest variability-index level, while fifteen-minute data-collection and schedule-development intervals yielded the most profitable schedules (or close to the most profitable) under the other variability-index levels. For the twenty-hour day, the fifteen-minute data-collection and schedule-development

intervals dominated across all variability-index levels.

Exhibit 5 illustrates mean schedule profitability by length of operating day, arrival rate, and variability-index level. The most interesting aspect of these results is that 5-minute data-collection and schedule-development intervals yielded the most-profitable schedules when averaged across all variability-index levels under the higher customer-arrival rate for the 14-hour operating day. Averaged across all variability-index levels, data-collection and schedule-development intervals of 15 minutes yielded the most-profitable schedules for the low customer-arrival rate for the 14-hour day, and at both arrival rates under the 20-hour day.

Synthesis and Implications

The results show that both high arrival rates and high arrival-rate variability work to reduce the ideal length of planning intervals. Let us consider why that happens. First, with high arrival rates, the

Exhibit 4: Profitability Results by Length of Operating Day

Day Length (hours)

copious data points on customer arrivals smooth the fluctuations in arrivals to give an estimate that is closer to its actual value, thus improving the accuracy of the projected customer-arrival rates. Those more accurate estimates of expected future demand allow relatively short data-collection intervals to outperform long data-collection intervals, even when there is relatively little arrival-rate variability. Second, with high arrival-rate variability, relatively long data-collection intervals fail to give an adequate picture of the

within-day variability, for the reason that variations within each period are not measured. As a consequence, the resultant labor-scheduling model does not offer a good match of capacity to demand. For example, if the average arrival rates in two consecutive fifteen-minute periods are ten customers and twenty customers, using thirty-minute data-collection periods would give an arrival rate of fifteen customers (the mean for the period). A restaurant in this instance would be relatively overstaffed in the first fifteen-minute

Exhibit 5:
Profitability Results by Length of Operating Day and Customer-Arrival Rate

period and relatively understaffed in the second fifteen-minute period. Thus, short data-collection intervals (and schedule-development intervals) are more useful than longer data-collection intervals (and schedule-development intervals) under higher arrival-rate variability.

Perhaps the most surprising finding in this study is the superiority of the shortest data-collection and schedule-development intervals. I have never seen five-minute periods examined in the literature, yet those periods were the best for high customer-arrival rates during a fourteen-hour day, when averaged across all levels of arrival-rate variability. Although I was unable to obtain solutions to the twenty-hour-day problems using five-minute schedule-development intervals (as explained in Exhibit 3), I see no reason why five-minute data-collection and schedule-development intervals would not also be superior to longer duration intervals for high arrival rates during a twenty-hour day.

Another argument for using short-duration planning intervals relates to real-time control. My simulation assumed that the employees performed as scheduled and that the schedule was not changed in real time. In actual operation, customer service can usually be improved and labor costs reduced by judiciously adjusting the schedule as the day unfolds. Real-time control includes altering the times at which breaks are taken, sending employees home early if customer demand is lighter than anticipated, and adding help if customer demand is higher than anticipated by calling scheduled employees in to work early, calling unscheduled employees in to work, or extending the length of some shifts. With short-duration planning intervals, a schedule contains more variety in the times at which employees start and finish work and start and fin-

ish breaks (as I have explained in a previous *Cornell Quarterly*).¹³ Schedule diversity of this kind could prove useful in improving the effectiveness of real-time control and reducing its cost. As such, scheduling using short-duration periods is likely to be preferable if a manager is considering real-time control.

Implications for Hospitality Managers

Clearly, a simulation tailored to a specific scheduling environment would offer the most information regarding the effectiveness of various duration-planning intervals. In lieu of such an investigation, one can measure the customer-arrival rate and the variability of that rate and compare those measures with the variability-index levels and the arrival rates used in this article. Even simpler, one could adopt fifteen-minute data-collection and schedule-development intervals because those intervals provided robust overall performance. If you are presently using data-collection or schedule-development intervals that are longer than fifteen minutes, you definitely need to review why you are doing so. Most likely you can improve the quality (i.e., profitability) of your schedules by adopting planning periods of fifteen minutes.

As explained here, it is often desirable to use periods shorter than fifteen minutes. However, using short-duration data-collection and schedule-development intervals is problematic at present, given the lack of scheduling applications that can accommodate those short intervals. I expect this article to spur the development of procedures to eliminate this deficiency. Hospitality managers can stimulate development of more effective planning intervals by asking the scheduling-solution providers to

support the use of short-duration periods. Unfortunately, the complexity of the calculations increases dramatically with short-duration periods (as may be seen from Exhibit 2), and so it is not as simple as just turning a switch and making it happen. In particular, scheduling tools have a tendency to "blow up" when they are given more flexibility than they were designed for, meaning that they no longer can find schedules in a reasonable amount of time. Thus, vendors of labor-scheduling systems may have to refine and rework their scheduling algorithms to accommodate the request for reduced planning-period durations.

Implications for Labor-scheduling Research

Baker suggested that one use planning intervals that vary in duration across the operating day, applying longer duration periods when the customer-arrival rate is relatively constant and shorter duration periods when arrival rates fluctuate.¹⁴ My findings here offer support for the potential value of such an approach, but its full benefit has yet to be established. My results also show that planning-interval duration should perhaps be adjusted based on the arrival rate itself and not just the arrival-rate variability.

A beneficial avenue for future investigations concerns the effectiveness of real-time control for improving customer service and lowering the labor costs. The case can be made that short-duration planning periods prove desirable even when the variability in customer-arrival rates is low, because of the flexibility inherent in the expanded diversity of the start and finish times for shifts and breaks.¹⁵

Given my finding regarding the value of short planning intervals, I see a need for scheduling procedures that can deal with

such brief intervals. Particularly troublesome is the literature devoted to the development of weekly (tour) schedules, which has used hour-long planning intervals almost without exception. The calculation-complexity problems posed by short-duration intervals, particularly in the tour-scheduling situation, pose interesting challenges to researchers. For the sake of the hospitality and other service industries, it is time for researchers and software developers to accept the challenge.

Endnotes

1. G. M. Thompson, "Labor Scheduling, Part 1: Forecasting Demand," *Cornell Hotel and Restaurant Administration Quarterly*, vol. 39, no. 5 (October 1998), pp. 22–31.
2. G. M. 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.
3. G. M. Thompson, "Labor Scheduling, Part 3: Developing a Workforce Schedule," *Cornell Hotel and Restaurant Administration Quarterly*, vol. 40, no. 1 (February 1999), pp. 86–96.
4. G. M. Thompson, "Labor Scheduling, Part 4: Controlling Workforce Schedules in Real Time," *Cornell Hotel and Restaurant Administration Quarterly*, vol. 40, no. 3 (June 1999), pp. 85–96.
5. G. M. Thompson and M. E. Pullman, "Scheduling Reliefs in Workforce Scheduling: In Advance, In Real-Time, Or Both," presented at the 26th annual meeting of the Decision Sciences Institute, Boston, November 19–21, 1995.
6. G. M. Thompson, "Accounting for the Multi-Period Impact of Service When Determining Employee Requirements for Labor Scheduling," *Journal of Operations Management*, vol. 11, no. 3 (September 1993), pp. 269–87.
7. Thompson and Pullman, loc. cit.
8. G. M. Thompson, "Labor Scheduling Using NPV Estimates of the Marginal Benefit of Additional Labor Capacity," *Journal of Opera-*

- tions Management, vol. 13, no. 1 (July 1995), pp. 67–86.
9. A. Brooke, D. Kendrick, and A. Meeraus, *GAMS, Release 2.25, A User's Guide* (San Francisco: Scientific Press, 1992).
 10. *Optimization Subroutine Library, Release 2* (Kingston, NY: IBM Corporation, 1991).
 11. G. M. Thompson, "Improved Implicit Optimal Modeling of the Labor-shift-scheduling Problem," *Management Science*, vol. 41, no. 4 (April 1995), pp. 595–607.
 12. Thompson, pp. 67–86.
 13. G. M. Thompson, "Controlling Action Times in Daily Workforce Schedules," *Cornell Hotel and Restaurant Administration Quarterly*, vol. 37, no. 2 (April 1996), pp. 82–96.
 14. K. R. Baker, "Workforce Allocation in Cyclical Scheduling Problems: A Survey," *Operational Research Quarterly*, vol. 27, no. 1 (1976), pp. 155–67.
 15. Thompson, pp. 82–96.

Gary M. Thompson, Ph.D., is an associate professor of operations management and executive director of the Center for Hospitality Research at the Cornell University School of Hotel Administration (gmt1@cornell.edu).