Labor Scheduling:

A Commentary

Labor scheduling—which involves putting the proper number of people in the right jobs at the correct time—is simple to say, but a challenge to accomplish.

BY GARY M. THOMPSON

In the late 1990s the *Cornell Hotel and Restaurant Administration Quarterly* published my four-article series on labor scheduling.¹ This commentary is intended for a wider audience than that of the earlier articles.

The four articles, though they can be considered as a primer to labor scheduling, were chiefly targeted at technically com-

petent people who have some degree of familiarity with labor scheduling. By contrast, this article is written in a way that, I hope, is accessible to those who are not technical experts and who are not particularly familiar with labor scheduling. In particular, missing in the four-article series was a thorough explanation of why a hospitality manager should care about the fine points of labor scheduling. I attempt to provide the motivation in this paper.

In this article I first define what I mean by labor scheduling and then tell you why I think it's important to master scheduling techniques. I then synthesize my earlier four articles, distilling and describing the major component tasks of labor scheduling. After that, I describe the characteristics I

¹See: G.M. Thompson, "Labor Scheduling, Part 1: Forecasting Demand," *Cornell Hotel and Restaurant Administration Quarterly*, Vol. 39, No. 5 (October 1998), pp. 22–31; 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; 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; and 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.

see as being important to look for when selecting a labor-scheduling system and talk about what I see as the potential pitfalls and major opportunities of labor scheduling. I close with some personal observations about trends in labor scheduling.

Prior to writing the four-paper series, I had developed a “black box” scheduler that is currently used by a prominent hospitality company for scheduling tens of thousands of employees.² Since writing the series, I have served as a consultant to a start-up company offering a web-based labor-scheduling solution. For the past six years I have been developing specialized labor-scheduling software for educational institutions (and selling that software commercially for over four years). These experiences have given me a deeper appreciation for the true challenges of labor scheduling.

What Is Labor Scheduling

Simply put, labor scheduling involves putting the right people on the right jobs at the right times (with regard to customer demand). As I stated in the first paper in the series, labor scheduling has the following four component tasks. They are:

...forecasting customer demand, translating those demand forecasts into employee requirements, scheduling employees, and fine-tuning the schedule in real time.

The first task is to predict customer demand for your service. This initial step involves forecasting characteristics of the service transaction that change over time, such as customer-arrival rates.

The second task is to calculate the number of employee hours required to satisfy demand predicted in step one. In other words, step two requires setting the number and skill levels of employees needed to serve customers adequately during some time period.

The third task is to develop the actual work schedule by taking into account employees' skills, desires, and requests, and then deciding who will do what work at what time.

The final task involves changing the work arrangement as required by actual demand. This final step ensures effective customer service.³

Why You Should Care about Labor Scheduling

As I see it, managers have three primary reasons to care about labor scheduling, starting with employee preferences. Employees generally have distinct preferences regarding their job, including the tasks to which they are assigned, when and how long their breaks are, with whom they take breaks, with whom they are working, the time of day they work, which days off they have, and whether their days off are consecutive. These preferences are commonly complementary across employees. By “complementary preferences,” I mean that employees differ in the characteristics they want in a schedule and they differ in the importance they place on those characteristics. By taking into account employees' preferences, and exploiting the complementary nature of those preferences, you'll get a work schedule that comes reasonably close to meeting employees' desires, which presumably translates to better on-the-job performance and better customer service.

The second reason to care about labor scheduling relates to the actual time spent on developing a labor schedule. Some firms use the “photocopier” method of developing a labor schedule, which entails creating this week's schedule simply by photocopying last week's schedule. Although this is a time-efficient approach, the schedules it yields are rarely any good, because the photocopier method does not adapt to changes in customer demand or to changes in employee availability or preferences.

Another method of developing the schedule is for a manager to build it on a weekly basis, drawing on his or her experience about what will

² This scheduler was integrated with a front-end graphical user interface provided by another vendor.

³ Thompson (Part 1), p. 23.

be necessary to serve customer demand and his or her knowledge about the preferences and availability of his or her employees. Though this method may actually yield good schedules, when a manager is conscientious and thoughtful it can require a substantial time commitment. Unfortunately, the time a manager spends developing a schedule leaves him or her less time for actually managing the employees and interacting with customers. Yet another approach would be to automate the schedule-development process, though this requires a computer system that can both forecast demand and account for all the employee idiosyncrasies with which adept managers have to deal. These computer systems are the holy grail of labor scheduling.

Profitability and effectiveness are the third reason to care about labor scheduling. Because a good labor schedule has the right people working the right jobs at the right times, good labor schedules deploy labor in the most effective manner. Effectively deploying labor translates to higher profitability, because short-term overstaffing and long-term understaffing are reduced. The resulting better, more consistent customer service translates to more future business. A related benefit is that good labor scheduling allows upper management to monitor performance more closely, both within and across units. Consistency in labor deployment across units is particularly important in chains, because it is one of the main reasons why service quality varies across units.

Major Component Tasks

Exhibit 1, shown on the next two pages, lists the major component tasks in labor scheduling. It also categorizes the tasks based on the primary task to which the task belongs and the frequency with which the task should be repeated.

Selecting a Workforce-scheduling System

Hospitality firms have many choices for workforce-scheduling systems.⁴ Windows-based workforce-scheduling software can be purchased

for as little as several hundred dollars, or firms can invest in customized systems costing hundreds of thousands of dollars. I believe that good scheduling systems have the following characteristics.

An intuitive graphical interface. The interface should facilitate tinkering with or editing a schedule. Although good systems will generate schedules that require a minimum of adjustment,

Managers should pay attention to scheduling to improve employee morale, save the managers' time, and boost the operation's profitability and service effectiveness.

systems rarely incorporate all relevant factors, and managers usually must make changes.

A good scheduling engine. The scheduling engine—the algorithms that actually develop the schedule—is harder to evaluate than the graphical interface because one cannot examine the engine directly. (I suggest a remedy for this problem below.) A good scheduling engine will incorporate effective logic. It should use a cross-period paradigm—considering the interactions of staffing decisions across planning periods—rather than a single-period paradigm. It should be holistic, meaning that it develops a schedule while considering employee availability, instead of shift-based, meaning that a schedule is first developed without regard to employee availability and then employees are matched to shifts. The engine should operate quickly and provide schedules that need a minimum amount of tinkering.

Employee preferences. A good scheduling system will require that employees identify their work preferences in advance, including ranking those preferences or identifying trade-offs among their preferences. The best system will exploit complementary preferences when they exist. By explicitly considering each employee's preferences, a good scheduling system will deliver the best possible schedule in terms of matching the number of employees scheduled to the ideal number of employees needed while at the same

⁴ Adapted from: Thompson (Part 3), *op. cit.*

EXHIBIT 1

Significant component tasks in labor scheduling

Primary task	Sub task	Frequency
Forecast customer demand ¹	Determine the nature of the work (controllable or uncontrollable). Controllable work is work for which there is some degree of timing control over when it can be performed. It typically can proceed without having customers present. Examples of controllable work are prep work in restaurants and housekeeping in hotels. Uncontrollable work is work over which there is little if any timing control, since it must be done when customers are in the system. Examples include serving customers in restaurants and checking-in guests in hotels.	Infrequent
	Identify those factors that generate the work; they are the key labor drivers. An example of a key labor driver for restaurant waitstaff would be the number of parties to be served.	Infrequent
	Determine whether the key labor drivers are time variant, meaning they vary over short time periods, or time invariant, meaning they are relatively stable over short time periods.	Infrequent
	Determine the time interval for tracking the time-variant labor drivers. In other words, this task identifies whether you should track the labor drivers using periods of 60 minutes, 30 minutes, 15 minutes, or using periods of some other duration. In general, time intervals of 15 minutes or shorter tend to work best. ²	Infrequent
	Develop forecasts of the time-variant labor drivers, using historical data. These forecasts can be modified, if desired, based on managers' special knowledge of future events.	Weekly
	Reduce the random variation in the forecasts of the time-variant labor drivers by smoothing the forecasts across periods. This step is important because the goal is to staff to the true level of customer demand and not staff to random variation in demand.	Weekly
	Measure the accuracy of the forecasts of the time-variant labor drivers. Forecast accuracy affects the number of employees needed to serve customers—less-accurate forecasts mean that you need a larger staff.	Weekly
	Determine the time period in which the controllable work can be performed. Some controllable work is itself dependent on the forecasts of time-variant labor drivers. For example, prep work in a restaurant kitchen, which is controllable work, must be done within time limits that are, in part, imposed by when customers will arrive.	Weekly

¹ For more details, see: Thompson (Part 1), *op. cit.*
² G.M. Thompson, "Planning Interval Duration in Labor Shift Scheduling," *Cornell Hotel and Restaurant Administration Quarterly*, forthcoming.

Exhibit continues on the next page

time satisfying employee preferences as much as possible.

Constraints. Hospitality businesses should be aware of their particular rigid and soft constraints. Rigid (or hard) constraints are those, like government regulations, that must be observed. Soft constraints, by contrast, are those, like employee preferences, that should be observed when possible. Knowledge of the operation's constraints will allow managers to evaluate the degree of congruence between a scheduling system's rigid and soft constraints and those identified by the hospitality firm. How a scheduling system deals with conflicting constraints can make the difference between an average system and a good one. Once again, the system's schedule should not require much managerial tinkering to get it into a usable form.

Flexibility. Good scheduling systems allow a high degree of scheduling flexibility, since flexibility yields better schedules. Good systems should allow various times for controllable work, multiple shift lengths, and a variety of break times. The system should also allow employees to be cross-scheduled within and across shifts. However, high flexibility greatly increases the complexity of a scheduling scenario and, consequently, may noticeably lengthen the time a system requires to develop a schedule.

Costs and benefits. Finally, when evaluating workforce-scheduling systems, one should not neglect the costs and benefits, both monetary and operational, of various systems. Inexpensive, off-the-shelf systems, for instance, will generally require that you fit your business operations to the system. Typically, this means that these systems will soon be abandoned. In contrast, a slightly more expensive, customized scheduling solution will fit the system to your business. Your firm's resources and potential benefits of a system will determine the best approach.

Potential Pitfalls

As with any type of management tool, there are potential pitfalls with labor-scheduling systems. The following are several to keep on your radar:

- *Using only aggregate measures of performance.* This is perhaps the biggest mistake

one can make related to labor scheduling. An example of an aggregate measure of performance would be labor dollars as a percentage of sales (that is, labor-cost percentage). What's missing from this performance measure is what's happening on a period-by-period basis within a week. If my target is for labor to be 20 percent of sales, I could hit that target by having labor be 30 percent of sales in half of my periods and 10 percent of sales in the other half of the periods (assuming sales were similar across periods). Thus, half the time I would have 50-percent more staff on hand than I needed, and the other half of the time I would have 50-percent fewer staff members than needed. A much better measure of performance would be the percentage of weekly periods in which labor is between 18 and 22 percent of sales.

- *Insufficiently testing the effectiveness of an autoscheduler.* If you're going to use a labor-scheduling system, then you'll want it to deliver the best possible schedule. Because users do not interact directly with the autoscheduler (the "black box" part of the system that actually develops the schedule), it can be difficult to evaluate the performance—a problem that I mentioned above. The best way to evaluate a part of the system that you cannot see is to create some logical test scenarios. The scenarios can be created in a way where it is obvious (to a perceptive manager) what the best schedule should be. These scenarios can then be fed to the scheduling system to see whether it can identify the best possible schedule. If a system cannot find the best schedule even for a simple scenario containing only a few employees, it is unlikely to perform well when it has to deal with real situations containing perhaps hundreds or thousands of employees. Simply put, if the autoscheduler is not up to snuff, then it's likely that the system won't be used, because the schedules won't be good enough. You might

EXHIBIT 1 (CONTINUED)

Primary task	Sub task	Frequency
Translate customer demand into employee requirements ³	Select a labor standard. The choices are productivity standards (which deliver a consistent level of labor utilization), service standards (which deliver a consistent level of service), and (my preference) economic standards (which deliver a consistent level of financial performance).	Infrequent
	Determine the number of persons, of each skill level, needed for each time interval in the week. Do this using the selected labor standard, the forecasts of time-variant labor drivers, and information on employee productivity and the accuracy of the forecasts.	Weekly
Build the schedule ⁴	Select a scheduling framework. The scheduling framework, of which several are common, ⁵ is the paradigm used to represent the labor-scheduling problems.	Infrequent
	Construct a labor schedule. Do this using information on employee availability, skills, and preferences, and the information on the number of employees needed in each period (see the accompanying text for a discussion of the characteristics of good work-force scheduling systems).	Weekly
Modify the schedule in real time ⁶	Monitor the schedule in real time, looking for capacity-demand imbalances. A capacity-demand imbalance exists when either more or less labor is available than is needed at a particular moment. The number of employees needed at a particular point would be determined based on the chosen labor standard, the current actual customer demand, and the productivity of the currently available employees.	Real time
	When a capacity-demand imbalance is detected, make a judgment as to whether it is likely to be short lived or long lived.	Real time
	If the capacity-demand imbalance is judged to be short lived, take actions like extending shifts, sending employees on breaks, or recalling employees from breaks.	Real time
	If the capacity-demand imbalance is judged to be long lived, take actions like calling additional employees in to work or sending employees home early.	Real time

³ For more details, see: Thompson (Part 2), *op. cit.*

⁴ For more details, see Thompson (Part 3), *op. cit.*

⁵ See, for example: G.B. Dantzig, "A Comment on Edie's 'Traffic Delays at Toll Booths,'" *Operations Research*, Vol. 2, No. 3 (1954), pp. 339–341; E.G. Keith, "Operator Scheduling," *AIIE Transactions*, Vol. 11, No. 1 (1979), pp. 37–41; 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, "Labor Staffing and Scheduling Models for Controlling Service Levels," *Naval Research Logistics*, Vol. 44 (1997), pp. 719–740.

⁶ For more details, see: Thompson (Part 4), *op. cit.*

think that the need for a test of this kind is obvious, but I have found that not one of the over 40 universities using my educational-labor scheduler has done any testing of this sort with my system.

- *Not using the system.* If you pay for a good system and it meets your needs, then use it. This may entail spending additional money for training. Investing in training may initially seem hard to justify if you

A key opportunity inherent in solid scheduling is in using employees in different functions, as needed.

have high turnover, but you can certainly assign the scheduling task to job positions that have less turnover. You might also take actions to reduce turnover. On the other hand, if you've paid for a good system and done the training, but your system is still not being used, perhaps it isn't as good as you thought. In that case, listen to the users of the system in your organization and then pass that information along to the system vendor. If the vendor doesn't listen, find one who will. From my own experience, the suggestions of users have been invaluable in improving every aspect of my educational-labor scheduler software.

- *Overblown claims of system vendors.* This pitfall is particularly insidious. Take a look at the website or promotional materials of any vendor of a labor-scheduling system and you're likely to see inaccurate claims, particularly regarding their autoschedulers (or optimization engines). In a quick review of some vendor websites, for example, I found one that claims to be able to "optimize the end-to-end labor scheduling process" and another that claims to offer "optimized schedules." Optimal, for those who don't know,

means "the best possible." In reality, given the state of today's computing power and the complexity of labor-scheduling problems, the systems typically are not able to find "the best" schedules. However, good systems will find good schedules and may, on rare occasions, actually stumble upon the best schedule.

Opportunities Afforded by Labor-scheduling Systems

In addition to the three items I discussed at the outset regarding why you should care about labor scheduling, I see two other opportunities. To me, one of the key opportunities is in cross-utilizing employees in different functions at different times. This is difficult to do if schedules are built manually, when the person constructing the schedule may be familiar only with the areas under his or her direct control. However, labor-scheduling systems can see the "big picture" and so deploy people in different jobs at different times. This capability, which exists today in labor-scheduling systems, means less understaffing and less overstaffing—in other words more effective labor deployment.

A second key opportunity involves enhancing the way that employee preferences are treated in labor-scheduling systems. All of the commercial labor-scheduling systems I've come across and all of the academic literature on labor scheduling have dealt with employee preferences from a hierarchical perspective. What this means is that these systems operate as though employees develop a hierarchy of what is important to them. However, people often make subtle distinctions between alternatives. The marketing literature on discrete-choice modeling offers a better alternative to the hierarchical preferences.⁵ Discrete-choice modeling involves presenting consumers two or more bundles of products and services that vary on different attributes and asking respondents to select their preferred choice. Along the

⁵ See: Rohit Verma and Gerhardt Plaschka, "The Art and Science of Customer Choice Modeling: Reflections, Advances, and Managerial Implications," elsewhere in this *Cornell Quarterly*.

same line, a labor-scheduling system that presents employees with choices between different possible schedules, and then uses that knowledge about their choices when developing a schedule, would take the employee side of labor scheduling to the next level of sophistication and performance.

Emerging Trends

The following are trends that seem apparent in the labor-scheduling domain.

- *Availability of on-line labor-scheduling systems.* As with a number of business functions, the internet has changed the labor-scheduling landscape. Several vendors now offer web-enabled labor-scheduling solutions. However, web enabling should supplement the key features of labor scheduling systems that I identified earlier and not replace those features. I make this observation because this has not always been the case.
- *More integrated management solutions.* Vendors of labor-scheduling systems appear to be taking a more holistic approach to labor scheduling. This translates into systems that tackle all four of the major tasks related to scheduling and, for example, integrating the labor-

scheduling system into time and attendance systems.

- *Labor outsourcing.* Another area that seems to be receiving more attention is labor outsourcing. If a hotel outsourced its housekeeping functions, for example, it would not have to worry about scheduling those housekeepers. In considering outsourcing, one should apply standard processes for evaluating decisions. In addition, I believe it can be helpful to think of employees as value providers and to ask how they can provide enhanced value, rather than think of the staff as a cost to be minimized. If you are confident that you are deploying your labor as effectively as is possible, only then are you ready to evaluate outsourcing.

Closing Thoughts

Much of this article contains opinions. Everyone has opinions, and yours may differ from mine. If you agree or disagree with my positions about labor scheduling, don't be afraid to share your ideas. You can share them directly with me via e-mail at gmt1@cornell.edu, or you can share them with the broader hospitality community on the discussion board on the CHR website (TheCenterForHospitalityResearch.org). ■

Gary M. Thompson, Ph.D., is an associate professor of operations management at the Cornell University School of Hotel Administration (gmt1@cornell.edu).