A Market-utility Approach to Scheduling Employees

By accounting for customers' reactions to service levels, a manager can schedule the optimum number of employees—both in terms of service levels and contribution to profit.

BY JOHN C. GOODALE, ROHIT VERMA, AND MADELEINE E. PULLMAN

Scheduling front-line service providers is a constant challenge for hospitality managers, given the inevitable tradeoff between service standards and operating expense. Traditional employee scheduling typically applies a cost-minimization approach to specify the level of front-line service providers who will be available to meet periodic demand. That cost includes the opportunity cost of lost customers,¹ which is part of the pseudo-costs of understaffing.² A con-

founding and often ignored effect, however, is the benefit generated by maintaining high service levels in a system where capacity exceeds demand.³ That is, scheduling more front-line service providers than the minimum level necessary to provide acceptable customer service (what might be considered to be overstaffing in some rubrics) may mean that customers receive service that is better than they expected (or what company standards prescribe).

In this paper we report on a scheduling approach that explicitly considers the interrelationships among customer preferences, customer demand, waiting times, and scheduling

¹ For example, L.L. Bean forecasted the number of callers that will abandon its catalogue-order system prior to talking to an agent, based on the level of staffing. After subtracting the projected number of those that will call back, Bean calculated a penalty cost due to understaffing. For a detailed explanation, see: B.H. Andrews and H.L. Parsons, "L.L. Bean Chooses a Telephone Agent Scheduling System," *Interfaces*, Vol. 19, No. 6 (1989), pp. 1-9.

² G.M. Thompson, "Labor Scheduling, Part 3," *Cornell Hotel and Restaurant Administration Quarterly*, Vol. 40, No. 1 (1999), pp. 86-96.

³ G.M. Thompson, "Labor Scheduling Using NPV Estimates of the Marginal Benefit of Additional Labor Capacity," *Journal of Operations Management*, Vol. 13, No. 1 (1995), pp. 67-86.

Market utility-based labor- and capacity-planning framework

Source: J.C. Goodale, R. Verma, and M.E. Pullman, "A Market Utility-based Model for Capacity Scheduling in Mass Services," *Production and Operations Management*, forthcoming (2003).

decisions. This approach, which we call the "market-utility model for scheduling" (MUMS), helps managers consider the dynamics of scheduling service employees. First, we discuss the components that make up this approach, which includes methods from customer-preferences modeling,⁴ service-capacity planning, and the four tasks of labor scheduling proposed by Thompson.⁵ Next, we'll show how the model applies to balancing queue lengths and operating costs for an airport food-court vendor. Finally, we discuss the value of MUMS for hospitality managers.

MUMS the Word

The MUMS framework comprises three primary components, namely, supply, demand, and economic consequences (see Exhibit 1).⁶ The supply component specifies the relationship between managerial decisions regarding service attributes and how they are reflected to the market, and how they affect the service design (structure). The service-attribute levels reflect customer preferences, capacity, and operating costs. The demand component captures the effect on the market of the service attributes from the demand component and forecasts customer-arrival rates based on the market utility.⁷ These customer-arrival rates are useful in projecting sales, but they are also projected inputs for the service-delivery system. The economic-consequence component accounts for the revenue generated by the market's projected response, and accounts for the costs of the managers' service-structure decisions. Below, we explain each of these components in detail.

Supply. The MUMS supply component contains the decision-making element of the model. Decisions regarding the service's physical structure, attributes, and the capacity level drive the

operation's expenses and appearance (and, ultimately, revenue). For example, staffing levels and employee schedules are keys in determining the service structure and the customers' anticipated waiting time. Staffing decisions of this kind are made in the context of the firm's competitive priorities,⁸ which are, in turn, developed with in-

Scheduling more front-line service providers than the minimum level necessary to provide acceptable customer service may mean that customers receive service that is better than they expected.

formation from marketing analysts regarding customers' preferences for service attributes and how much value customers place on waiting time as a service attribute. This relationship is reflected in Exhibit 1 by the arrow pointing from marketing toward the managerial-decisions box, indicating that this information is collected and synthesized to make decisions regarding service capacity. Conversely, service attributes borne out of managerial decisions regarding competitive priorities are promoted to potential customers via marketing (as depicted by the arrow pointing from managerial decisions to marketing, and from marketing to the demand component).

The managerial-decisions element of Exhibit 1 also contains the function that specifies expenditures on the service structure (shown by the arrow pointing toward the operations box). For example, the employee schedule is part of the specification of the front-line structure for serving customers. The arrow pointing from operations to the economic component reflects costs

⁴R. Verma, G. Plaschka, and J.J. Louviere, "Understanding Customer Choices: A Key to Successful Service Management," *Cornell Hotel and Restaurant Administration Quarterly*, Vol. 43, No. 4 (December 2002), pp. 15–24.

⁵Thompson, *op. cit.*

⁶See: J.C. Goodale, R. Verma, and M.E. Pullman, "A Market Utility-based Model for Capacity Scheduling in Mass Services," *Production and Operations Management*, forthcoming (2003).

⁷Verma, Plaschka, and Louviere, *op. cit.*

⁸For an explanation on competitive priorities in services, see: A.V. Roth and M. Van der Velde, "Operations as Marketing: A Competitive Service Strategy," *Journal of Operations Management*, Vol. 10, No. 3 (1991), pp. 303–328.

for the service structure (for example, corresponding labor costs). However, the decisions that specify the structure use information from operations (represented by the arrow to managerial decisions) regarding existing structural strengths and weaknesses.

The Demand Component. Three elements make up the demand component: market size, market share, and customer-arrival rate. For any given market size, based on customer preferences, the relative importance (or utilities)

Too many customers may degrade service levels to the business's ultimate detriment.

of various service attributes can be estimated using the choice-modeling procedures described by Verma, Plaschka, and Louviere.⁹ This information can be used to estimate market share for each competing firm. Then, the actual customer-arrival rate for a specific facility is calculated by multiplying market size with market share (which is a function of the set of utilities of service attributes as perceived by the customers). A fundamental insight regarding the relationship between the framework's supply and demand components is that the level of customer arrivals affects service-attribute levels (e.g., expected waiting time, perception of quality). Although being busy might generate more revenue in the short term, it may not be good for customer service (with probable negative effects on future revenue). In short, changes to a service system that generates long waits (or in general, low-quality customer service) can ultimately reduce future market share.

This fundamental insight (i.e., that too many customers may degrade service levels to the business's ultimate detriment) has two important ramifications on market utility-based scheduling of service capacity. First, a schedule is generated based on projected periodic

employee requirements (usually 60- or 30-minute planning periods) that reflect the levels of service attributes (for example, expected queue time and variety of services available).¹⁰ However, because we generally cannot treat each period independently (given that employees do not work just hour-long or 30-minute shifts), we risk over- or understaffing in each individual period. Thus in any individual period, customers may experience service levels that are better or worse than hoped for by management, and this will affect service attributes experienced by customers (in this case, waiting time). The second important ramification is that changes in the levels of any service attribute may cause corresponding changes in market share (and thus, customer arrivals), which will then change determinant attributes related to waiting (that is, the service operation will be busier and customers' wait times may be longer).

Economic Component. The framework's economic component integrates the service structure from the supply component and the customers' decisions in the demand component by specifying operating profit. Variable costs, such as materials (for example, food to be sold in restaurants) can be combined with the revenue per customer multiplied by market share to yield the contribution toward fixed costs from operations. Direct server costs (or employee costs), however, which are a function of labor and capacity decisions, are separated, for our purposes. The contribution from operations (which includes employee costs) minus fixed costs yields operating profit.

Scheduling at an Airport Food-court Restaurant

To demonstrate MUMS, we analyzed a fast-food restaurant that operates in an international terminal at a major airport in the United States. The restaurant—a McDonald's—is one of four retail food-service operations located in the terminal, where each is open between 6:00 AM and

⁹ Verma, Plaschka, and Louviere, *op. cit.*

¹⁰ For an analysis of queue-related factors and management options, see: Carolyn U. Lambert and Thomas P. Cullen, "Balancing Service and Cost through Queuing Analysis," *Cornell Hotel and Restaurant Administration Quarterly*, Vol. 28, No. 2 (August 1987), pp. 68-72.

Choice set #11	Restaurant 1	Restaurant 2	Restaurant 3	Restaurant 4	None
Brand name	McDonald's	Local restaurant	La Prefeda/ Goya Products	Subway/ Boston Market	
Variety	Burger, fries, ice cream	Pizza, lasagna, pasta, salads, soups	Hot dogs, fries, tamales, nachos, burritos, tacos, enchiladas	Sandwich, soup, ice cream, udan noodle soup, salads	
Wait time (before ordering, in minutes)	5-6	0-2	3-4	0-2	
Service time (minutes)	0-2	3-4	5-6	3-4	
Language	English	English, Spanish, Japanese	English, Spanish	English, Spanish	
Menu-picture display	Yes	No	Yes	No	
Price (meal + drinks)	\$4	\$4	\$10	\$7	
I would purchase food from:					

9:00 PM daily. The McDonald's that we studied was the largest of the four restaurants in this food court. This McDonald's is, of course, part of a large, international chain that specializes in hamburgers, but also offers a limited number of other menu items. The structure of this service operation is also familiar: multi-server and single phase. That is, customers exit upon completion of their transaction with a particular front-line service provider. Each customer's order is handled by a single server, with several servers on duty at any one time. Waiting lines can form in front of each server at the counter, but customers are free to switch lines at will (and do), and thus we assume that there is, in effect, a central waiting area rather than multiple waiting lines. The managerial decision that we addressed is the scheduling of the restaurant's front-line service employees.

We began our analysis by examining preferences of customers using discrete-choice analysis (DCA) methods.¹¹ DCA can model customer utility (preferences) for a service in response to

experimentally designed profiles of service attributes. Recent studies demonstrated that market-utility models developed from carefully conducted DCA experiments can accurately predict market share for various types of products and services.¹²

DCA experiments involve careful design of service profiles (that is, specific service attributes for each restaurant) and choice sets (that is, groups of attributes for a particular restaurant) in which two or more service alternatives are offered to decision makers in the form of different sets of attributes. Customers are asked to evaluate the options and choose one (or none). Based on the experimental design, the customers' choices are a function of the attributes of each alternative, respondents' personal characteristics, and unobserved effects (a random factor). To develop the market-preference structure for the

¹¹ As described by: Verma, Plaschka, and Louviere, *op. cit.*

¹² For guidelines on designing and conducting DCA studies for restaurants, see: R. Verma, G.M. Thompson, and J.J. Louviere, "Configuring Service Operations in Accordance with Customer Needs and Preferences," *Journal of Service Research*, Vol. 1, No. 3 (1999), pp. 262-274.

Food-court arrival rate

airport-terminal restaurant, we performed the following steps: (1) identification of attributes, (2) specification of attribute levels, (3) experimental design, (4) presentation of alternatives to respondents, and (5) estimation of choice model.¹³ Experimental attributes were brand name, menu language, menu variety, picture display, price, service time, and waiting time. Twelve choice sets were created, each containing four descriptions of a restaurant based on various high, low, or missing levels of the experimental attributes. Five hundred randomly selected customers were asked to choose one option from each choice set (or choose none). A sample choice set is shown in Exhibit 2 (on the previous page).

Exhibits 3, 4, and 5 show the arrival-rate levels and how those levels were translated into contribution from operations. First, with the sales information from the restaurant's point-of-sale system and its baseline market-share projections, we estimated daily demand for the entire food court, broken down into hourly planning periods (see Exhibit 3). With only a slight resemblance to typical demand patterns for regular restaurants where demand peaks during meal time, the food court's demand pattern was largely a function of airlines' arrivals and departures.

Second, the restaurant captured various amounts of this market depending on the attributes of its service offering. Exhibit 4 shows how waiting in line affected market share, by presenting the arrival rates as a function of staffing level. We developed an algorithm that tied waiting-time projections to market-share estimates based on a decision-support system, which used the customer-choice models that we developed.¹⁴ The market-share estimate could include the design configurations for any combination of the seven attributes and generate an estimated market share for each vendor. For this study, the seven design attributes for each food vendor were initially set to their actual configuration. To examine the market's sensitivity to waiting in line, we changed only the attribute of waiting in the

¹³ As recommended by: *Ibid.*

¹⁴ M. Pullman, R. Verma, and J.C. Goodale, "Service Design and Operations Strategy Formulation in Multicultural Markets," *Journal of Operations Management*, Vol. 19, No. 2 (2001), pp. 239-254.

line before service at the restaurant under study, while configurations in the market-share estimate were kept constant for the other restaurants in the food court. Combining the market-share estimate (arrival rate) with the waiting-line model completed a circular relationship (that is, expected waiting time from the waiting-line model was an input for the market-share model). The arrival rate converged to an equilibrium point that represented the market share that managers would achieve for various staffing levels.¹⁵ Large chunks of market share were gained by adding employees at 8:00 AM. However, each additional employee added beyond three increased the marginal arrival rate at a substantially smaller pace than what is seen by adding the first three.

Given that declining arrival-rate increment, the restaurant eventually reached a point at which the cost of an additional employee would be more than the incremental increase in revenue from serving a few more customers. Multiplying the customer-arrival rates from Exhibit 4 by the operations contribution per customer and subtracting the incremental-employee cost provided the incremental operations contribution in a particular planning period. For illustration, Exhibit 5 shows the incremental operations contribution net of employee costs, assuming a marginal operations contribution (excluding employee costs) of \$4 per customer and a wage rate (employee costs) of \$6 per hour. The threshold at which additional employees provide a net benefit was considered across the length of their shift (that is, one must evaluate the cumulative net benefit over the entire shift). For this illustration, we used only part-time employees who were available for four-hour shifts. However, any shift-length combination can be considered.¹⁶

The final stage in implementing the framework was to create an optimal solution for the restaurant's shift-scheduling problem. For the optimal solution, we generated the numbers of employees that should be working each shift for the four-hour shift patterns.¹⁶ The optimal staff-

EXHIBIT 4

Customer-arrival rate as a function of staffing levels

EXHIBIT 5

Operations contribution per incremental employee

¹⁵ The details can be found in Goodale *et al.* (forthcoming, 2003).

¹⁶ Generated according to: *Ibid.*

EXHIBIT 6

Hourly staffing levels to maximize operating contribution

No. of Employees	Time														
	0600	0700	0800	0900	1000	1100	1200	1300	1400	1500	1600	1700	1800	1900	2000
1															
2															
3															
4															
5															
6															
7															
8															
9															
10															

Note: Shaded area indicates number of employees on duty for each hour.

EXHIBIT 7

Optimum employee-scheduling pattern

Employee number	Time														
	0600	0700	0800	0900	1000	1100	1200	1300	1400	1500	1600	1700	1800	1900	2000
1															
2															
3															
4															
5															
6															
7															
8															
9															
10															
11															
12															
13															
14															
15															
16															
17															
18															
19															
20															
21															

Note: Shaded area indicates an employee's shift.

ing levels would provide the distribution of employees shown in Exhibit 6. In a four-step format, the shift scheduling procedure was as follows:

- (1) Define the objective as maximizing the contribution from scheduling shifts, such that each shift supplies the sum of the periodic contributions (from Exhibit 5) in their working periods at the new level of employees.
- (2) Formulate the problem, such that an array of binary variables (i.e., 0 or 1) is generated and can be superimposed on a look-up table (such as one containing the actual numbers that Exhibit 5 is reflecting). This array might ideally look like Exhibit 6, except it will have 0s for open blocks and 1s instead of shaded blocks.
- (3) Evaluate the sum of the products from multiplying corresponding elements of the superimposed arrays. The problem can be formulated as a non-linear, integer, or linear-programming (LP) problem and solved accordingly.¹⁷
- (4) Implement the optimum schedule of 21 employees as shown in Exhibit 7.¹⁸

The Value of MUMS

The MUMS approach has value for hospitality managers at two levels. At a strategic level, MUMS provides a conceptual framework of an important, dynamic relationship that is often difficult to grasp. That is, managers attempt to align operational structure with competitive priorities to gain market share. Managers' success (or lack thereof) will be reflected by the increase (or decrease) in arrivals to the operation. In the case where managers successfully capture additional market share by changing service attributes, they must have sufficient capacity to provide desired levels of customer service, or poor performance in customer service (waiting in line) may cause market share and arrivals to decrease.

¹⁷ For the LP formulation of this integer programming problem, such that the simplex method in Solver can be used to solve the problem optimally and instantly, see: *Ibid*.

¹⁸ Note that if employee wages were entirely uniform, then back-to-back four-hour shifts in this schedule might be converted to eight-hour shifts (with perhaps a flexible lunch break), which would reduce the total number of employees by approximately 25 percent. Alternately, the problem can be solved with a variety of shift lengths.

The challenge for the manager is to find the "best" level of any attribute such that the restaurant's service structure will support the resulting customer volume. Of course, a manager will not attempt to do this at any cost, but so that operating contribution is maximized. This is the role of MUMS. The restaurant illustration in the paper demonstrated how managerial decisions regarding scheduling will align service structure with service offerings.

At an operational level, MUMS provides a framework for managers to evaluate the economic effects of their decisions regarding service offerings (or attributes) and service capacity (the number of employees to schedule). In a dynamic sense, the MUMS approach links service attributes to employee scheduling, and employee scheduling to market share and contribution margin.

A number of innovative approaches for matching service supply and demand have emerged, and all of these methods have merit. Aside from effective employee scheduling, hospitality operations can: (1) make more use of employees by cross-training them, (2) bring in part-time employees for peak customer volumes, (3) manage demand levels with variable pricing, and (4) inform or train customers about the service so that they increase their participation delivering that service. However, even when employing these other methods the fundamental dynamism of setting staffing levels remains. MUMS accounts for these dynamic relationships in a contribution-maximizing manner.

As a final note, we argue that treating operations as a competitive or strategic priority precludes a cost-minimizing orientation. As such, scheduling methods should embrace tools from functional areas that assist in using market information, for example, DCA. Used as a strategic framework or an operational tool, the MUMS framework offers exciting opportunities for gaining insight into service operations and developing performance-enhancing methods. ■

John C. Goodale, Ph.D. (top photo), is an assistant professor at the Charles H. Lundquist College of Business at the University of Oregon (jgoodale@LCBmail.uoregon.edu). **Rohit Verma**, Ph.D. (middle photo), is an associate professor at the David Eccles School of Business at the University of Utah (rohit.verma@business.utah.edu). **Madeleine E. Pullman**, Ph.D. (bottom photo), is an assistant professor at the College of Business at Colorado State University (mellie.pullman@mail.biz.colostate.edu).