

1-1-2003

Dedicated or Combinable? A Simulation to Determine Optimal Restaurant Table Configuration

Gary M. Thompson Ph.D.
Cornell University, gmt1@cornell.edu

Follow this and additional works at: <http://scholarship.sha.cornell.edu/chrpubs>

 Part of the [Hospitality Administration and Management Commons](#)

Recommended Citation

Thompson, G. M. (2003). Dedicated or combinable? A simulation to determine optimal restaurant table configuration [*Electronic article*]. *Cornell Hospitality Report*, 3, 4-16.

This Article is brought to you for free and open access by the The Center for Hospitality Research (CHR) at The Scholarly Commons. It has been accepted for inclusion in Center for Hospitality Research Publications by an authorized administrator of The Scholarly Commons. For more information, please contact hlmdigital@cornell.edu.

Dedicated or Combinable? A Simulation to Determine Optimal Restaurant Table Configuration

Abstract

Using a computer simulation, one can determine what the optimum table arrangement would be for restaurants of various sizes that accept walk-in customers only and take no reservations. At issue is whether the restaurateur can gain more revenue when its tables are dedicated to seating parties of specific sizes (for example, parties of one and two people would be served at 2-tops, while parties of one to four people would be served at 4-tops) or whether the restaurant should use tables that can be combined as needed according to party size. The simulation predicted that combinable tables would prove most useful in a small restaurant with a small average party size. Combining tables in that situation increased revenue per available seat hour by about 2 percent compared to having only dedicated tables. In a large restaurant or any restaurant with a large average party size, the simulation found that dedicated tables were superior to combinable tables. A loss in productivity occurs when some number of tables are held out of service until adjacent tables become available (so that the tables can be combined to seat a large party). The simulation found that the most efficient approach is for a restaurant's table-size mix to match its customer party size mix, since doing so increases the restaurant's effective customer-service capacity. However, that customer mix cannot always be known before a restaurant is constructed, and that mix might change during different dayparts. Moreover, the simulation makes certain assumptions that may need further examination, and it does not take into account such aesthetic factors as customers' reactions to a particular restaurant layout.

Keywords

table arrangement, party size, revenue per available seat hour RevPASH, Tablemix

Disciplines

Business | Hospitality Administration and Management

Comments

Required Publisher Statement


© [Cornell University](http://www.cornell.edu). This report may not be reproduced or distributed without the express permission of the publisher

CORNELL

CHR Reports

Dedicated or Combinable? A Simulation to Determine Optimal Restaurant Table Configuration

by Gary M. Thompson, Ph.D.



The Center for Hospitality Research
AT THE SCHOOL OF HOTEL ADMINISTRATION



CHR Reports

is produced for the benefit of the hospitality industry by
The Center for Hospitality Research at Cornell University

Cathy A. Enz, Executive Director
Glenn Withiam, Director of Publications Services

Price US \$50.00

Copyright © 2003 by Cornell University

Dedicated or Combinable? A Simulation to Determine Optimal Restaurant Table Configuration

by Gary M. Thompson, Ph.D.

Executive Summary

Using a computer simulation, one can determine what the optimum table arrangement would be for restaurants of various sizes that accept walk-in customers only and take no reservations. At issue is whether the restaurateur can gain more revenue when its tables are dedicated to seating parties of specific sizes (for example, parties of one and two people would be served at 2-tops, while parties of one to four people would be served at 4-tops) or whether the restaurant should use tables that can be combined as needed according to party size. The simulation predicted that combinable tables would prove most useful in a small restaurant with a small average party size. Combining tables in that situation increased revenue per available seat hour by about 2 percent compared to having only dedicated tables. In a large restaurant or any restaurant with a large average party size, the simulation found that dedicated tables were superior to combinable tables. A loss in productivity occurs when some number of tables are held out of service until adjacent tables become available (so that the tables can be combined to seat a large party). The simulation found that the most efficient approach is for a restaurant's table-size mix to match its customer party-size mix, since doing so increases the restaurant's effective customer-service capacity. However, that customer mix cannot always be known before a restaurant is constructed, and that mix might change during different dayparts. Moreover, the simulation makes certain assumptions that may need further examination, and it does not take into account such aesthetic factors as customers' reactions to a particular restaurant layout.

Dedicated or Combinable? A Simulation to Determine Optimal Restaurant Table Configuration

by Gary M. Thompson, Ph.D.

THIS STUDY USES A COMPUTER SIMULATION to model the seating patterns and table arrangements of full-service restaurants that accept only walk-in customers. Rather than accept reservations, restaurants of this type either seat guests as they arrive or (when the restaurant is busy) ask their guests to wait in a queue. Such chains are common in the United States; Cracker Barrel, Red Lobster, and TGI Friday's are examples of this type of restaurant. When a table is available, a host or hostess seats the parties separately (that is, each party sits at its own table, even if that means empty seats).

The goal of this simulation is to predict the optimum table arrangement in terms of revenue for restaurants of this type. Specifically, this study examines whether such restaurants should be configured with tables of various sizes that can be dedicated to a particular party size or else configured with small tables (say, deuces and 4-tops) that can be combined as needed to seat large parties. To illustrate the difference in the two types of configuration, consider a restaurant that serves parties of one through eight people. Using dedicated tables, one might use a mix of 2-, 4-, 6-, and 8-top tables, where parties of one and two patrons are served at the 2-tops, parties of three and four guests are served at the 4-tops, parties of five and six people are served at the 6-tops, and parties of seven and eight diners are served at the 8-tops. On the

other hand, if tables can be combined, the restaurant might be composed entirely of deuces that are pushed together as needed.

One can make arguments in favor of either configuration. One would expect, for example, that a restaurant would gain efficiency from the added flexibility of configurable tables. Then again, that restaurant would lose efficiency when a table must be placed on hold until it can be combined with an adjacent table that is still in use. A restaurant with dedicated tables generally would not require the idle, on-hold time of configurable tables, but that restaurant might lose efficiency if the "wrong size" party shows up when the appropriate table is not available. That is, if six people show up but only 4-tops are available, the restaurant is in an awkward position, just as would occur if two people found themselves seated at an 8-top, when that's the only table available.

I measure performance in this paper based on RevPASH—revenue per available seat hour—for each table arrangement.¹ The following are my goals for this investigation:

- Determine which strategy gives stronger revenues, dedicated or configurable tables;
- Determine the extent of RevPASH differences between dedicated and configurable tables; and
- Examine how the ideal mix of tables differs under dedicated and configurable table designs.

To answer these questions, I created a simulation to model each type of restaurant. After a discussion of the relevant literature, I introduce the restaurant-simulation model, describe the experiment, present results, and offer conclusions.

A Revenue-management Approach

Table management of the kind I discuss here fits within the conceptual framework of revenue management, since its focus is on maximizing the revenue that a restaurant can achieve in a given time period. Restaurant revenue management has only recently begun to receive attention, most notably from Sheryl Kimes and her coauthors, who have introduced the time-based revenue performance measure—RevPASH—that I use in this

¹ See: S.E. Kimes, R. B. Chase, S. Choi, E. N. Ngongzi, and P. Y. Lee, “Restaurant Revenue Management,” *Cornell Hotel and Restaurant Administration Quarterly*, Vol. 40, No. 3 (June 1998), pp. 40–45.

paper.² Kimes *et al.* have presented strategies for restaurant revenue management³ and suggested steps for implementing revenue management in restaurants.⁴

Beyond the Kimes work, the literature on restaurant-table optimization is scant. A recent unpublished paper Kimes and I wrote discussed our analysis of the ideal table mix for a particular mid-scale, full-service restaurant.⁵ When we allowed tables

In some cases, combining tables may allow a restaurant to serve up to 30-percent more customers.

to be combined, we found that the ideal table mix would enable the restaurant to process approximately 30-percent more customers, without increasing wait times, than would the restaurant’s existing (dedicated) table mix.

² *Ibid.*

³ S.E. Kimes, D. I. Barrash, and J. E. Alexander, “Developing a Restaurant Revenue-management Strategy,” *Cornell Hotel and Restaurant Administration Quarterly*, Vol. 34, No. 5 (October 1999), pp. 18–30.

⁴ S.E. Kimes, “Implementing Restaurant Revenue Management: A Five-step Approach,” *Cornell Hotel and Restaurant Administration Quarterly*, Vol. 34, No. 3 (June 1999), pp. 16–21.

⁵ S.E. Kimes and G.M. Thompson, “Restaurant Revenue Management at Chevy’s: Determining the Best Table Mix,” Working paper 07-05-02, School of Hotel Administration, Cornell University (www.hotelschool.cornell.edu/chr/research/working/chevysrevenue.pdf).

I am not aware of any other research that addresses the particular issue that I examine in this paper. The key aspect of table mixes in restaurants is the ability to combine resources—that is, tables—to serve larger parties than would otherwise be possible.

TABLEMIX Simulation Model

A restaurant-simulation model that I have developed was a key component of this research. This model, which I

The key aspect of table mixes in restaurants is the ability to combine resources—that is, tables—to serve larger parties than would otherwise be possible.

call TABLEMIX, simulates how customers use tables in a restaurant.⁶

TABLEMIX can be run in a graphical mode, in which case it will display the status of the restaurant—showing which tables are occupied, the number of seats occupied at each table, which tables are on hold, and the number, size, and waiting time of parties waiting for tables. On-hold tables are those that will be combined with an adjacent table to seat a larger party (for example, two adjacent 4-tops would be combined into an 8-top), when the party at the adjacent table completes dining.

⁶ TABLEMIX was developed with Microsoft's Visual Basic 6.0® and runs under Windows operating systems.

Assumptions. The assumptions for the TABLEMIX simulation are that the time between party arrivals is exponentially distributed;⁷ that the space occupied by a table is proportional to the number of its seats; and that parties do not share tables (so, for example, two separate two-person parties would not be seated together at a 4-top).

The inputs for TABLEMIX enable users of the tool to closely match the simulation with the characteristics of actual restaurants. The inputs include the length of the dining window to be simulated; the expected number of party arrivals, by 15-minute period for the daypart (the “dining window”); the number of days of “operation”; the probabilities of various-size parties; the maximum number of waiting parties; the distribution of dining duration (normal or log normal); and the table-assignment rule (that is, the hostess either assigns an available table to the largest possible party for that table or else to the party waiting the longest).

The simulation inputs that are specified for each size party include the mean and standard deviation of dining time by party size, the maximum wait the party will tolerate, and the party's revenue (or contribution value).

TABLEMIX can be used to evaluate a specific restaurant configuration,

⁷ That is, there is a high probability of a short time between party arrivals and a low probability of a long time between party arrivals, which is a common situation in practice.

or it can be used to search for the best restaurant configuration. When simulating a specific restaurant configuration, one must specify the number of tables, as well as the number of seats and the position of each table, and identify which tables can be combined with which other tables. If TABLEMIX is used to identify the best restaurant configuration, one must specify which size tables can be used along with the limit on the number of seats in the restaurant. In this case, TABLEMIX will enumerate and evaluate all possible table configurations that use the full complement of seats.⁸

The simulation's outputs include the average wait by party size, the number and value of parties served and lost, the number and value of customers served and lost, the utilization of each size table by 15-minute period, and seat utilization by 15-minute period.

Experimental Design

The experiment that I designed had one managerially controllable factor—that being the degree of combinability of tables (five levels)—and two environmental factors—namely, the number of restaurant seats (50 and 200) and the mean party size (2.5, 3.5, and 4.5). The probabilities I used for each size party are presented in Exhibit 1.

Since the largest party size was eight people, the largest tables had to be 8-tops for the restaurant with

⁸ Or as close as possible to the full complement of seats given the allowable table sizes.

EXHIBIT 1
Probability of different size parties, under the three levels of mean party size

Party Size	Mean Party Size		
	2.5	3.5	4.5
1	0.16	0.05	0.02
2	0.51	0.17	0.12
3	0.15	0.34	0.16
4	0.10	0.27	0.20
5	0.04	0.08	0.23
6	0.02	0.04	0.13
7	0.01	0.03	0.09
8	0.01	0.02	0.05

dedicated tables. Thus, for both size restaurants, I allowed the use of 2-, 4-, 6- and 8-tops. For the 50-seat restaurant, this results in a total of 185 distinct table mixes, all of which use the full complement of 50 seats.⁹ For the 200-seat restaurant, there is a staggering total of 8,037 separate table mixes that use the full complement of seats.¹⁰


As I indicated, combinability—the single managerially controllable factor in the experiment—had five levels, varying from no combinability (i.e., dedicated tables) to high combinability (i.e., many small tables). I measure combinability as follows: measure the number of table pairs that can be

⁹ For a 50-seat restaurant, one mix would be 25 two-tops and another would be 6 eight-tops plus a two-top.


¹⁰ For a 200-seat restaurant, one mix would be 100 two-tops, and another would be 25 eight-tops.

EXHIBIT 2


Examples of 50-seat restaurants composed of 25 two-tops, with differing degrees of table combinability




100% of tables can be combined



50% of tables can be combined



30% of tables can be combined




10% of tables can be combined

Note: Lines indicate places where tables can be combined.

combined as a percentage of the maximum number of table pairs that could be combined given no facility constraints. The five levels of combinability are: 100 percent, 50 percent, 30 percent, 10 percent, and zero. Exhibit 2 shows an example of the combinability levels for a 50-seat restaurant where all the tables are 2-tops. Tables that can be combined are shown by connecting lines.

Building a “restaurant.” When TABLEMIX evaluates one of the enumerated table arrangements, it randomly places the tables in rows in the restaurant (keeping the restaurant as close to square as possible). It then randomly makes adjacent pairs combinable, unless the desired proportion of combinable tables has been achieved (based on the table-combinability factor). For example, with any of the 8,037 distinct table mixes in the 200-seat restaurant, there can be an immense number of different configurations, based on which tables are placed adjacent to which other tables and which tables may be combined. By randomly placing each table and by randomly specifying which tables are combinable, I am, in effect, simply sampling one of the possible configurations with that mix of tables. As a way to reduce the effect of purely random variation that this sampling might introduce into the results, when I report which table mix worked best, I do so based on an average of the top 5 percent of the table mixes for each combinability level (which is nine table-mix configura-

EXHIBIT 3**Party arrival rates, by 15-minute period, for the 200-seat restaurant**

rations for the 50-seat restaurant and 402 configurations for the 200-seat restaurant). Compared to simply picking the best-performing table mix, averaging the results for the top 5 percent of table mixes enables me to get a better sense of the effects of table combinability.

Dinner hour. Assumptions that I made in the experiment were a 55-minute mean dining time for all parties; a log-normal distribution of dining times;¹¹ a \$10-per-person average check (dining value) for all parties; a maximum tolerable wait of 90 minutes for all parties; a table-assignment rule that gave the next available table to the largest waiting party; simulating 150 days of opera-

¹¹ A log-normal distribution of service times occurs quite commonly in restaurants. Log-normal distributions look much like a typical normal distribution (bell curve), except that one tail is longer. In this case, there is a higher probability of extended dining durations than would occur with a standard normal distribution.

tion (equivalent to 1.5 years' worth of operation, where there were two peak days per week); that no more than 10 parties could be waiting for the 50-seat restaurant or 40 parties for the 200-seat restaurant; and that tables would be combined only for party sizes of five and larger. I also assumed a peak unconstrained-demand level that would result in a seat utilization of 100 percent. Setting the demand at this level means that some parties will not be served, but it also ensures that any differences resulting from combinability will be apparent. Finally, I used a five-hour peak dining window, but measured RevPASH only after the first 90 minutes of operation, since the first 90 minutes represent the ramp up to the peak dining period. The party-arrival rates that I used for the 200-seat restaurant are shown in Exhibit 3.¹²

¹² The party-arrival rates for the 50-seat restaurant are one-quarter of those for the 200-seat restaurant.

EXHIBIT 4

Calculation of the naïve ideal number of tables required for a 50-seat restaurant and a mean party size of 2.5 people

Table size (Number of seats)	2	4	6	8
Tables required per party	0.67 ^a	0.25 ^b	0.06 ^c	0.02 ^d
Naïve ideal number of tables	11.713	4.371	1.049	0.350

Notes:

^aEqual to the sum of the probabilities of party sizes 1 and 2 (=0.16+0.51)

^bEqual to the sum of the probabilities of party sizes 3 and 4 (=0.15+0.10)

^cEqual to the sum of the probabilities of party sizes 5 and 6 (=0.04+0.02)

^dEqual to the sum of the probabilities of party sizes 7 and 8 (=0.01+0.01)

EXHIBIT 5

Naïve ideal number of tables for the three mean party sizes and the two restaurant sizes

Restaurant size	Mean party size	Table size (Number of seats)			
		2	4	6	8
50	2.5	11.713	4.371	1.049	0.350
50	3.5	2.750	7.625	1.500	0.625
50	4.5	1.400	3.600	3.600	1.400
200	2.5	46.853	17.483	4.196	1.399
200	3.5	11.000	30.500	6.000	2.500
200	4.5	5.600	14.400	14.400	5.600

A base table mix, which I call the naïve ideal table mix, can be determined using the following formula:

$$T_s = \frac{TPP_s}{\sum_{j \in S} j * TPP_j} * SEATCAP$$

where T = naïve ideal number of tables with seats; TPP_s = number of tables of s seats required per party; SEATCAP = and S = the set of allowable table sizes.


If dedicated tables are used, then the number of tables of s seats that are required per party (TPP) is given by:

$$TPP_s = \begin{cases} \sum_{j=1}^s \text{PROBParty}_j, & \text{if } s = \text{smallest allowable table size} \\ \sum_{j=\text{nxtsm}_s+1}^s \text{PROBParty}_j, & \text{otherwise} \end{cases}$$


where nxtsm_s = largest table size with fewer than s seats.

Exhibit 4 presents the calculation of the naïve ideal table sizes for the mean party size of 2.5 people and a 50-seat restaurant, and Exhibit 5 does the same for all three levels of mean party sizes and both restaurant sizes.

EXHIBIT 6
RevPASH, by party size, for the two restaurant sizes




50-seat restaurant



200-seat restaurant

EXHIBIT 7

RevPASH under the combinability levels, as a percentage of the RevPASH for the dedicated tables (0 combinability), averaged across all environmental experimental factors




Combinability versus Dedication

Exhibit 6 (on the previous page) graphs the RevPASH by mean party size. In the 50-seat restaurant when the party size was small (2.5) or medium (3.5), RevPASH increases as table combinability increases to the 50-percent level, but declines at 100-percent combinability. For largest mean party size in the 50-seat restaurant and for all party sizes in the 200-seat restaurant, RevPASH generally declines as the level of combinability increases. Exhibit 7 shows the average RevPASH for the combinability levels expressed as a percentage of the RevPASH for dedicated tables (i.e., no combinability). High combinability levels lower the relative RevPASH, with complete combinability yielding a 1.5-percent lower RevPASH than dedicated tables.


Party mix. The best table mixes are illustrated in Exhibit 8. There are four observations of interest to be made from this exhibit. First, the naïve ideal table mix is similar to the best table mix that occurs with dedicated tables (no combinability). This result was what I expected, since the naïve ideal table mix was calculated with the assumption of dedicated tables. Second, mean party size had a large effect on the best table mix. Third, higher combinability meant a greater number of 2-tops, a result particularly noticeable with the largest mean party size. Fourth, when the party size was small, the table mixes were similar across all combinability levels.

EXHIBIT 8


Naïve ideal table mixes and the best performing table mixes by combinability level, for the six combinations of restaurant size and mean party size




a) 50-seat restaurant, 2.5 mean party size




b) 200-seat restaurant, 2.5 mean party size




c) 50-seat restaurant, 3.5 mean party size



d) 200-seat restaurant, 3.5 mean party size



e) 50-seat restaurant, 4.5 mean party size



f) 200-seat restaurant, 4.5 mean party size

Note: Percentages indicate the proportion of tables that may be combined.

The Use of Dedicated Tables

First, I must address the question of why combinable tables worked best in the 50-seat restaurant with small and medium mean party sizes. With a 50-seat restaurant, seats in general are at a premium. When the mean party size is small or medium, the simulation indicates that having a dedicated 8-top is not worthwhile. There just are not enough large parties to occupy the seats on that big a table. At the same time, under the assumptions used in this simulation, if the largest table is a 6-top, any party larger than six people would be lost. Combinable tables in this case allow the restaurant to serve parties larger than six people without incurring the excess capacity that a dedicated 8-top would require.

Large restaurant, large parties.

The other question to examine is why dedicated tables worked better with all party sizes in the 200-seat restaurant and with large party sizes in the 50-seat restaurant. The short answer is that those cases had a sufficient number of large parties to justify at least one 8-top. The restaurant with an 8-top will not lose parties for want of a large enough table. In these cases, having combinable tables does not allow the restaurant to serve any additional customers. Indeed, allowing tables to be combined means that tables will sometimes be placed on hold, which interferes with productive capacity.

Two other factors can further reduce the value of combinable tables. First, I assumed that seats would be

conserved when tables are combined, even though that might not be the case in a real restaurant. In the simulation, for example, combining two 4-tops yields an 8-top. Some restaurants would lose seats in this situation, where combining two 4-tops might yield only a 6-top. Combinability is less attractive in that situation because of the loss of effective capacity that occurs when tables are combined. Another consideration is the time required to combine tables on the fly. The simulation assumed that no work time was required to combine tables, but that process obviously removes servers from their duties and would again diminish effective capacity. The net effect of the assumptions I just discussed is to favor combining tables in the simulation. Thus, the fact that my results showed only a slight advantage for combining tables and did so only in certain situations should give one serious pause before considering the use of combinable tables.

Getting the mix right. The importance of finding the table mix that is appropriate for one's restaurant cannot be overstated. My discussion so far is based on the assumption that a restaurateur has been able to identify (and install) the best mix of tables for the restaurant's customer mix. Another way to look at the results, though, is to consider a particular table mix and see whether dedicated or combinable tables work better. For the 50-seat restaurant, the 50-percent combinability level yielded higher RevPASH than did dedicated tables

in 332 of the combinations of a specific table mix and mean party size, and lower RevPASH in 222 instances. For the 200-seat restaurant, the 50-percent combinability level yielded higher RevPASH than did dedicated tables in 13,035 of the combinations of a specific table mix and mean party size and lower RevPASH in 11,074 iterations. These results suggest that, in general, if one does not have the best mix of dedicated tables in one's restaurant, it is probably advisable to install combinable tables.

Caveats. One should be wary of relying solely on this simulation, for two reasons. First, the analysis has not considered how customers may react to the aesthetics of a particular table arrangement. Its aim was simply to analyze restaurants as productive systems. Second, the best table mix will vary, as we have seen, with the size of the restaurant and the mean party size. Thus, an essential determinant of the most effective table mix is the customer mix. Although one never knows with certainty what the customer mix will be until a restaurant actually opens, chain-restaurant operators can make a reasonable assessment by applying party-size information from other restaurants in locations with similar demographics. The operators of independent restaurants will have a more difficult time making that assessment. In any case, it would seem prudent periodically to evaluate whether the table mix in one's restaurant is consistent with the restaurant's customer mix.

Issues for Future Analysis

This report is the first in what I expect to be a stream of research studies, since there are at least five issues raised in this paper that I would like to address. The first of these is the effect of the table-assignment rule on restaurant performance. This analysis used the rule that tables would be assigned to the largest waiting party, but I wonder what would happen if a

In some cases, combining tables may allow a restaurant to serve up to 30-percent more customers.

different rule were used. The largest-party rule can result in different levels of service for different size parties. For example, this rule means that a singleton would, on average, have to wait longer than parties of two, because when any table becomes available it will be given to the largest waiting party (i.e., the pair). Thus, singletons would be assigned to tables only when there were no waiting parties of two. That may not be realistic for actual restaurant operation.

The second issue is identifying which sizes of tables should be located adjacent to and combinable with which other tables. For example, is it better to make two 4-tops combinable, or is it better to combine a 6-top and a 2-top, or a 4-top and two 2-tops? The

answer to this may well depend on the restaurant's party-size distribution and perhaps the establishment's overall size.

The third issue to examine is the size of parties for which tables should be combined. In this research I combined tables only for parties of five and larger, but better performance may be achieved by allowing tables to be combined for smaller (or only for larger) parties. I suspect that this issue is likely related to the probability distribution of party sizes. For example, the best answer for a case like this experiment, where there were no parties larger than eight people, is likely to be different than in the case where the largest parties have than more eight people.

Fourth, I suspect that combining tables may be more desirable when the restaurant has a various peak demand periods that are characterized by different customer mixes (say, lunch and dinner, or after theater).

Finding a single, dedicated table mix that works well across different dining periods is likely to be harder, I think, than finding a single mix of combinable tables that works well in this case.

Finally, a colleague and I are currently working on faster ways of determining the best table mix. Making a complete enumeration of all possible table mixes becomes impractical when more than four table sizes can be used or the size of the restaurant exceeds about 200 seats. In the experiment described here, the 200-seat restaurant had a total of 8,037 different table mixes. However, if the simulation had allowed 10-top tables in the 200-seat restaurant, then there would have been 44,559 distinct table mixes—an increase of over 400 percent. Since simulating 8,037 table mix alternatives already takes over four hours of computer time, there is a definite need for a tool that can quickly find the best-performing table mix. □

Gary M. Thompson, Ph.D., is an associate professor of operations management at the Cornell University School of Hotel Administration (gmt1@cornell.edu). He holds a B.Sc. with First Class Honours from the University of New Brunswick, an M.B.A. from the University of Western Ontario, and a Ph.D. in operations management from Florida State University. His research interests are in the application of quantitative techniques to complex hospitality problems. His primer on workforce scheduling appeared in the Cornell Hotel and Restaurant Administration Quarterly. His research has also appeared in Decision Sciences, Journal of Operations Management, Journal of Service Research, Management Science, and Naval Research Logistics. He is also founder and CEO of SchedulExpert, Inc., a software firm focused on scheduling products.



About the Center for Hospitality Research at Cornell University

The Center for Hospitality Research is the premier source for in-depth research on the hospitality industry. The Center constitutes a community of researchers and industry sponsors working in concert to develop and dispense new knowledge that directly bears on the hospitality industry. Membership in the Center is open to any individual or corporation interested in furthering high quality research for the betterment of the hospitality industry.

The goal of the Center for Hospitality Research is to inform scholarship in hospitality with an industry perspective. Development of the CHR's research efforts are augmented by industry representatives' expertise. By enhancing industry access to meaningful research, the CHR creates a community of researchers and industry sponsors.



The Center for Hospitality Research

AT THE SCHOOL OF HOTEL ADMINISTRATION

hosp_research@cornell.edu or 607-255-9780



The Center for Hospitality Research
AT THE SCHOOL OF HOTEL ADMINISTRATION
