

Measurable

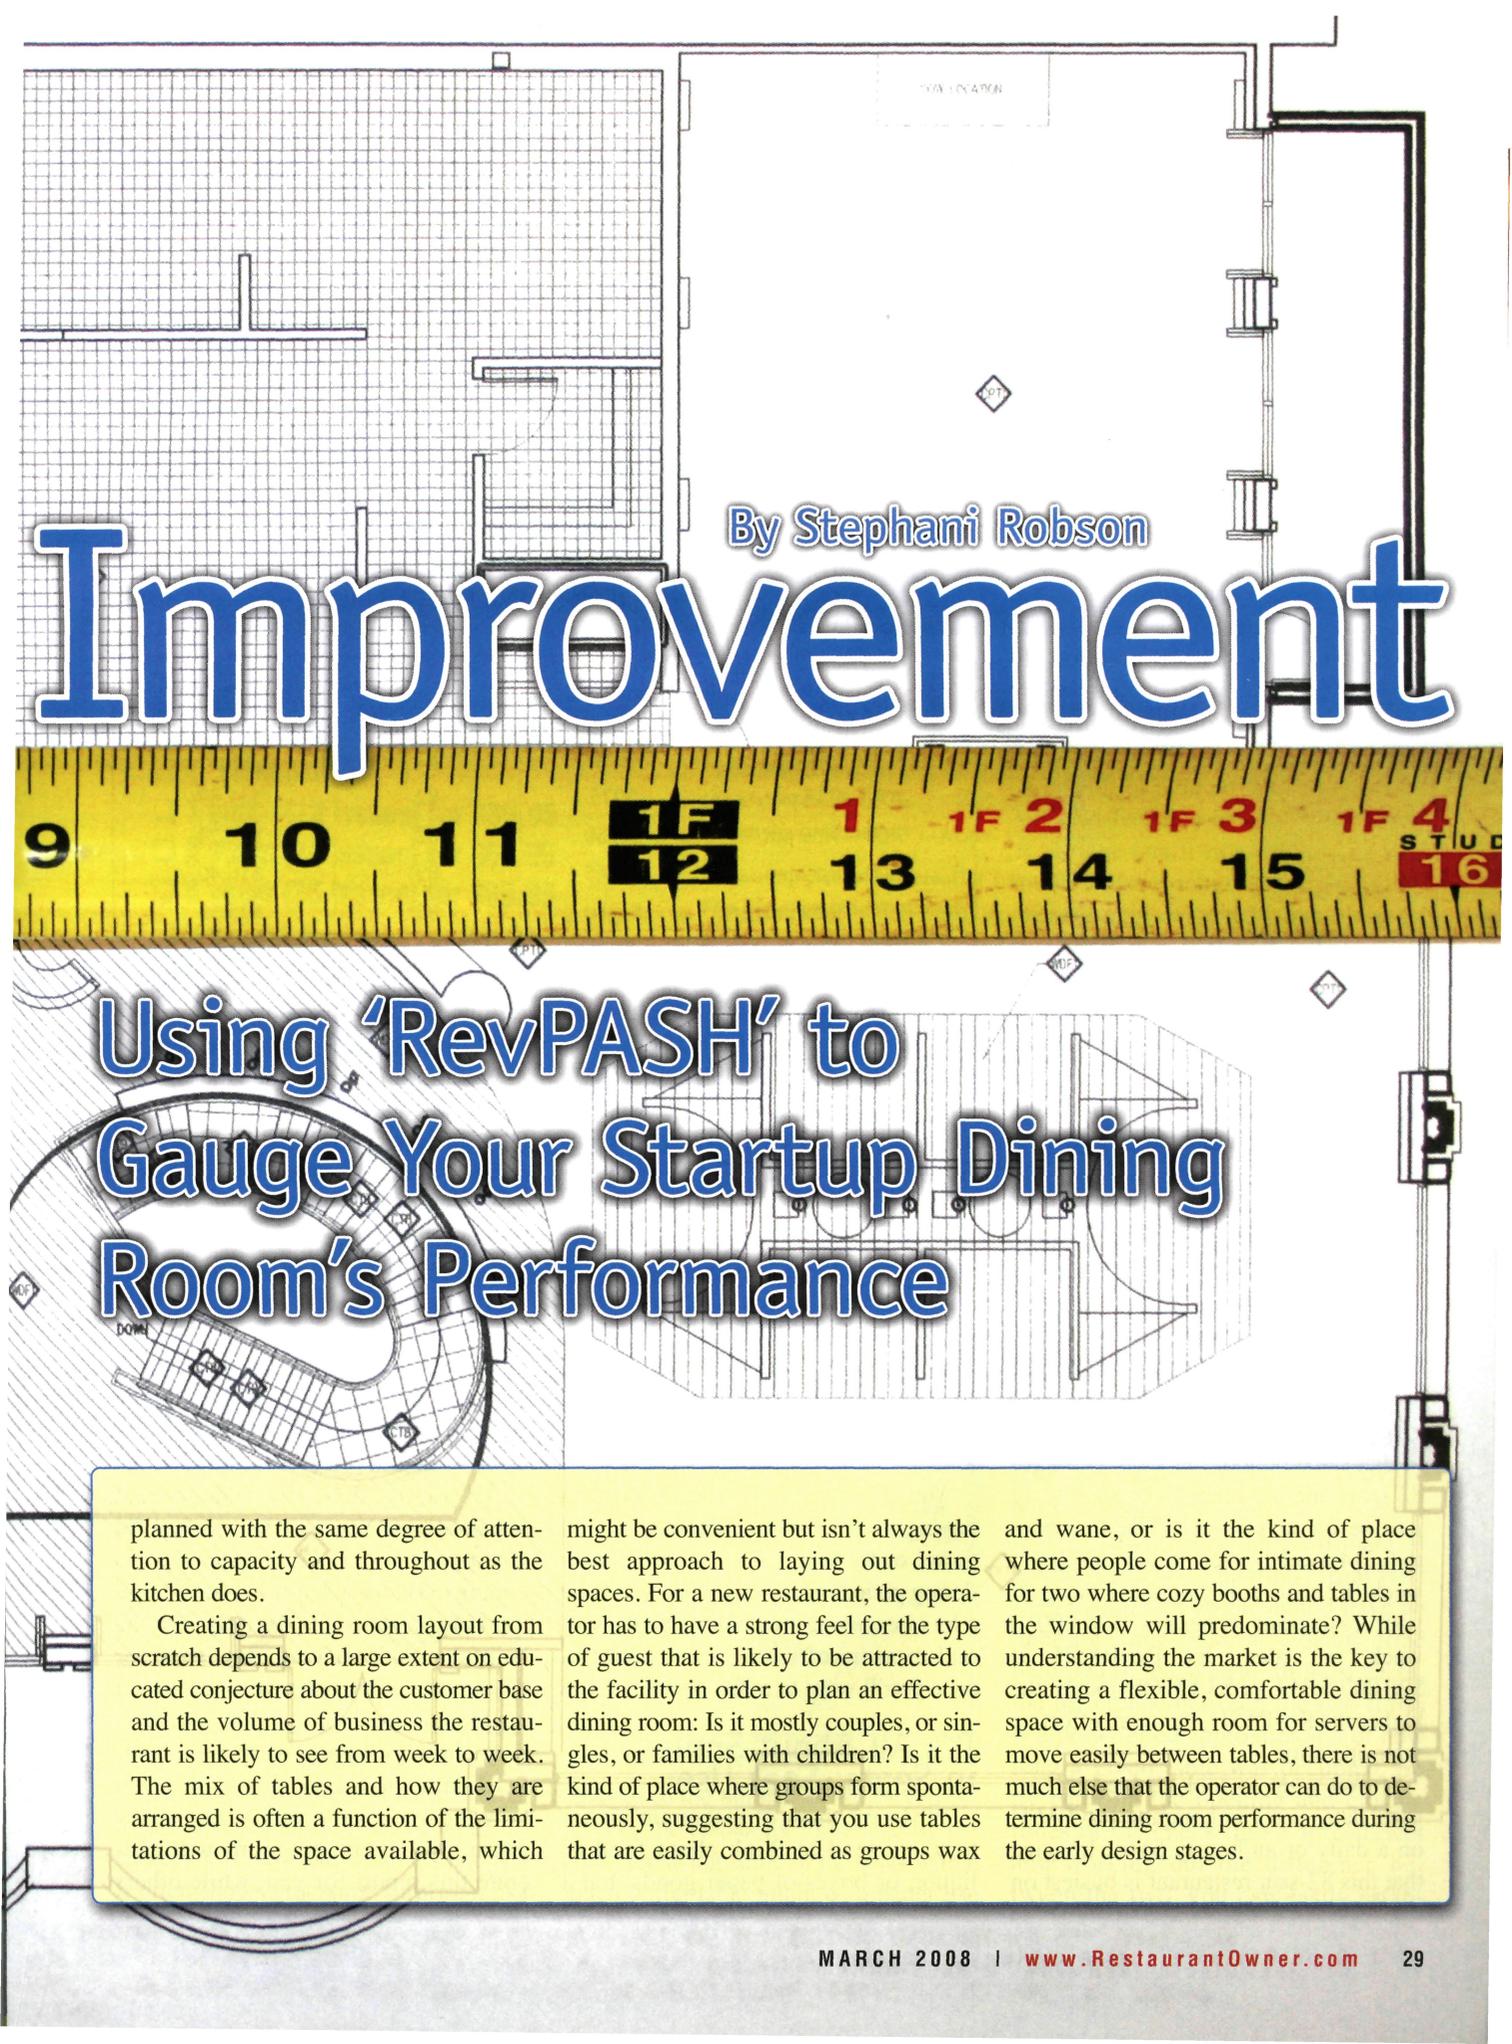
As a restaurant operator, you're probably used to getting reports. Last night's sales. The weekly liquor inventory. Last month's payroll. Some days, it may feel like you're drowning in paper and numbers, and that you're measuring and evaluating everything you possibly can in your operation. But are you measuring how well your dining room is working? No, not your servers, but the *design* of the

dining room itself? If you aren't, you could literally be leaving money on a lot of tables.

Design Matters

The design of your restaurant can be compared to the layout of a factory (although I admit that the analogy lacks a certain sense of romance). Factories are designed to maximize potential output while making the best possible use of

every square foot. Restaurants need to be designed the same way. Your output is satisfied customers, and a good restaurant design maximizes the number of these happy diners that you can "process" in a fixed period of time. On the kitchen side, this translates to having the right equipment in the right configuration, housed in the right amount of space. But the dining room is part of the "factory" too, and therefore needs to be



By Stephani Robson

Improvement

Using 'RevPASH' to Gauge Your Startup Dining Room's Performance

planned with the same degree of attention to capacity and throughout as the kitchen does.

Creating a dining room layout from scratch depends to a large extent on educated conjecture about the customer base and the volume of business the restaurant is likely to see from week to week. The mix of tables and how they are arranged is often a function of the limitations of the space available, which

might be convenient but isn't always the best approach to laying out dining spaces. For a new restaurant, the operator has to have a strong feel for the type of guest that is likely to be attracted to the facility in order to plan an effective dining room: Is it mostly couples, or singles, or families with children? Is it the kind of place where groups form spontaneously, suggesting that you use tables that are easily combined as groups wax

and wane, or is it the kind of place where people come for intimate dining for two where cozy booths and tables in the window will predominate? While understanding the market is the key to creating a flexible, comfortable dining space with enough room for servers to move easily between tables, there is not much else that the operator can do to determine dining room performance during the early design stages.

But when your restaurant is already operating, you have a wealth of information available to you that can allow you to adapt the dining room to be as efficient and effective as possible. Most of that information is right there in your point-of-sale (POS) system, just waiting to be put to work: table number, dining duration and party size. Together, these three pieces of information can provide a snapshot of your operation's front-of-house productivity as well as pinpoint areas where you could be doing better, sometimes with very little additional effort or investment. Unfortunately, POS systems don't give you this information directly; you have to ferret it out before you can use it. If you are comfortable with spreadsheet tools like Microsoft Excel, you may be able to download your POS data into a spreadsheet format so you can sort it and analyze it. (Ask your POS vendor whether this is possible and, if it is, to give you detailed written instructions on how to go about it. Some systems are much easier to extract data from than others.) In this article, I'll show you how each of these measures can help you have a better front-of-house operation and how you can compute these measures for yourself.

RevPASH — Better Than Sales Figures

The first measure of dining room performance is Revenue per Available Seat Hour, or RevPASH. RevPASH is easy to calculate: Just divide your revenue for a given period by the available seat hours in that same period. Let's say you have 82 seats and are open from noon to 11 p.m. seven days a week. That translates to 6,314 available seat hours (ASH) per week (82 seats x 11 hours x 7 days). If your restaurant generates \$20,000 a week in sales, its RevPASH is \$20,000/6,314, or \$3.17. Obviously, the higher the RevPASH, the more effective your operation.

If you want to look deeper at your sales data, you can compute RevPASH on a daily or an hourly basis. Imagine that this 82-seat restaurant is busiest on

Saturday nights when every seat is full between 7 and 10 p.m. Its hourly sales data might look like this:

<i>Time of Day</i>	<i>Total Sales</i>
5:00-5:59 pm	\$350
6:00- 6:59 pm	\$975
7:00 – 7:59 pm	\$1250
8:00 – 8:59 pm	\$1345
9:00 – 9:59 pm	\$1280
10:00 – 10:59 pm	\$775
Total	\$5975

Dividing sales by the number of seats in your restaurant for each hour gives your hourly RevPASH:

<i>Time of Day</i>	<i>Total Sales</i>	<i>Seats</i>	<i>RevPASH</i>
5:00-5:59 pm	\$350	82	\$4.27
6:00- 6:59 pm	\$975	82	\$11.89
7:00 – 7:59 pm	\$1250	82	\$15.24
8:00 – 8:59 pm	\$1345	82	\$16.40
9:00 – 9:59 pm	\$1280	82	\$15.61
10:00 – 10:59 pm	\$775	82	\$9.45

Periods with higher RevPASH indicate times when it's probably in your best interest to turn tables, assuming that there are patrons waiting, whereas when RevPASH is relatively low, upselling might be the best strategy for increasing revenue. So for the restaurant in our example, promoting desserts early and late in the evening makes sense, but perhaps you might want your servers to hold back on promoting desserts during the peak dining hours of 7-9 p.m. when getting your waiting patrons into your dining room is likely to result in higher revenues per hour.

You can increase RevPASH either by having your existing customers spend more, or by bringing in more business during periods of the day when you have a lot of excess seat inventory. So the second measure of dining room performance you'll want to make is seat use.

It's All About Smiles in Seats: Seat Use

When we think of restaurant inventory, we tend to think about food or liquor, or boxes of paper goods, but a

restaurant's real inventory is time: We have only so many seats, and only so many hours in a day that those seats can be occupied by paying customers. (I'm going to focus this discussion on full-service restaurants that do not have a lot of takeout volume.) And unlike most of the other kinds of inventory you keep, seat inventory is tremendously perishable: If a seat goes unoccupied for an hour, the potential revenue that this seat could be generating is lost forever. But taking stock of how well seat inventory is used — also termed "seat utilization" — is rarely done by independent restaurateurs. You may have a feel for when your restaurant is quiet (i.e., is underused) and when it is packed, but tracking seat use accurately can be helpful in identifying day parts or days of the week when your restaurant could be doing more, and the exercise may show you some surprises. Your POS data should be able to tell you how many covers you sold in any given hour. Some programs will generate this report for you, while others

will need to have the information transferred to a spreadsheet program like Excel so you can do the analysis. (This is a great project for that sharp summer intern you've taken on from the local college's hospitality management program.) The number of covers per hour, when compared with the number of seats in your restaurant, gives you a snapshot of seat use for that hour. If you are selling more covers in an hour than you have seats, that suggests that your dining room is highly productive. Few restaurants achieve this level of activity for extended periods; in fact, seat use of more than 75 percent in any given hour is considered to be very good dining room performance.

The line graph below is an example of a seat use analysis for a 100-seat bistro-style restaurant surrounded by office buildings. The proportion of total seats that are occupied are plotted hour by hour, with a different linetype indicating each day of the week. This chart tells us that this particular restaurant gets most of its demand at lunchtime, which isn't surprising given its location close to lots of lunch customers. But even during its busiest hours, the restaurant still has about half of its seats unfilled.



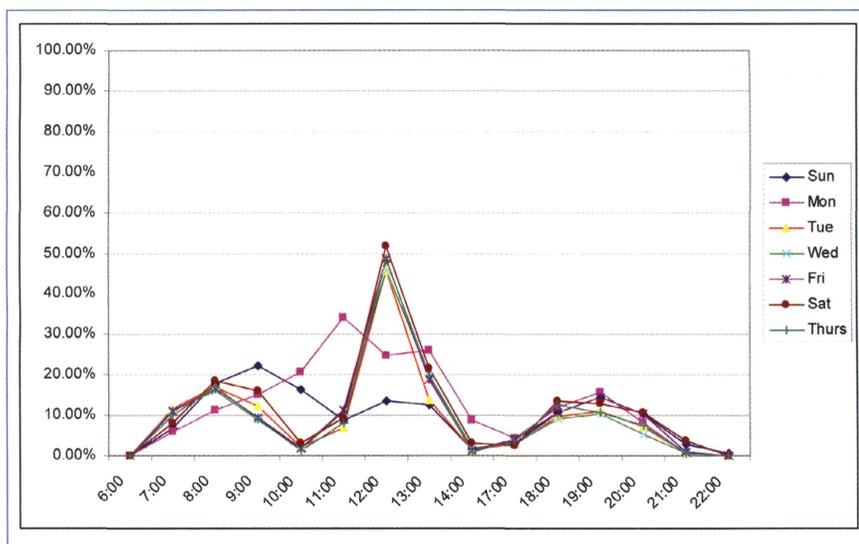
This restaurant uses mostly 'anchored' tables; that is, tables located next to some kind of architectural feature such as a window or a low wall, to make guests feel more comfortable. Photo by Robert Barker, Cornell University Photography.

the different lunch pattern on Mondays that you can see in the chart for our bistro example. Looking into the causes of these patterns and variations can help you plan your staffing, figure out when to offer promotions, or even suggest

of space and possibly attract more lunchtime patrons who would prefer the faster option of self-service during the limited lunch hour. Tracking seat use gives you the hard facts to help you make these kinds of decisions.

Should They Stay or Should They Go? Dining Duration

Another part of tracking dining room performance is knowing how long your seats are occupied during a typical meal. While no one wants to measure guests' dining experiences with a stopwatch, there is a simple way to deduce a reliable estimate of length of stay or "duration" from your POS data. Each check should have an opening time and a closing time recorded in the system. In most cases, the difference between the two times is a pretty good approximation of how long a given table was occupied. So by entering the opening and closing times of the check into a spreadsheet and working out the difference between the two, you can easily estab-



This analysis probably confirms what a good operator already knows: when the restaurant is busy and when it's not. But it often shows you some surprises, like

changes to the number of tables you offer. For the restaurant illustrated by the chart, replacing some of the seats with a lunch buffet table might be a better use

lish average durations for your different day parts. True, there are some parties that will linger long after their check has closed, and even if your staff is exceptionally attentive, there is likely to be a short lag from the time the party is seated to when the first order is entered into the POS. If it makes you feel more confident in your results, add a few minutes onto the length of stay of each party to account for these variations.

Here is a summary of average durations by day and meal period for the bistro restaurant illustrated in the previous section. You can see that Saturday's an unusual day, in that breakfast and lunch guests dine relatively quickly, whereas dinners take quite a bit longer on that night of the week. Compare this with Sundays when there is a much longer duration at breakfast; an extensive buffet brunch is offered that day — that's why there is no average duration for lunch on Sundays for this particular restaurant. With information like this, you can manage your reservations a bit more accurately, knowing to space them a bit further apart Thursday through Saturday than you would during the rest of the week. Or maybe you need another breakfast server on Mondays to help get guests through their meal quickly.

	<i>Monday</i>	<i>Tuesday</i>	<i>Wednesday</i>	<i>Thursday</i>	<i>Friday</i>	<i>Saturday</i>	<i>Sunday</i>
Breakfast	48 min.	43 min.	39 min.	37 min.	42 min.	38 min.	58 min.
Lunch	53 min.	54 min.	59 min.	51 min.	58 min.	44 min.	---
Dinner	72 min.	73 min.	63 min.	77 min.	71 min.	82 min.	59 min.

In addition to giving you a clearer picture of dining room performance, taking a look at dining duration can also help you identify problem spots in your service. Wide variation in dining duration can sometimes be a signal that something is up: a server requires more training (or just more help on the floor), or that the kitchen has trouble getting some items out the door. Duration is usually longer for larger parties, which is understandable given the greater demands that larger groups place on your staff. A good way to spot possible

problems is to compare length-of-stay estimates for similar party sizes and compute the average duration for each party size. For example, all your parties of two might take an average of 57 minutes to dine with you. If you see a number of two-tops that took a substantially longer time to dine, check whether these were all on the same day, or had the same server, or ordered the same menu items. When we have performed studies of restaurant data at the Cornell School of Hotel Administration, we've calculated party size averages and standard deviations (a measure of how much each check varies from the average) and find that if the standard deviation for duration is larger than the average duration, the restaurant may be having serious problems with timely service. Some tables may be taking twice the average time to get through the meal, which is quite possibly the result of uneven service. Doing this more detailed analysis at least raises the flag that one part of your service needs to be examined more closely. (Excel can calculate standard deviations from your list of durations for you, saving you from having to hire a mathematician.)

What's a Good Table?

But isn't it also possible that some diners just take a long time to enjoy their experience? We've all heard tales that diners at the "good" tables (the ones in the window, or that big corner booth) take longer over their meal than diners at less desirable tables; but is this restaurant folklore really true?

Multiple studies of restaurant POS data that look at duration by table location suggest that there *is* a difference, but not in the way that you might expect. Diners who sit at desirable tables like those in the window actually don't stay any longer than patrons at less coveted places in the dining room. In fact, the longest durations are often at the *least* desirable tables.

What's more, there appear to be differences in spending across table types in many restaurants. In some of our studies, we've found that desirable tables have often had higher average checks than undesirable tables do. In fact, you can combine the two results — duration and average check — into a single measure, which we call spending per minute (SPM). SPM is easy to calculate: Just divide the average check from each party by its duration. We've found that spending per minute tends to be higher at good tables as opposed to bad ones. "Good tables" are the ones that guests seem to request most often; in the window or in a booth (everybody seems to love booths), whereas "bad" tables are those in the middle of the dining room that aren't by a wall, partition or other structure. (Tables that are positioned next to permanent architectural features are called "anchored" tables, and guests of many types appear to prefer anchored tables in almost all cases, according to recent research. See the photograph on Page 21 for an illustration.) Bad tables tend to have more traffic around them because they have aisles on all sides of them, increasing the chances for patrons to be bumped, or just feel less comfortable and relaxed because of the activity close by.

Here's what we found when we looked at spending per minute across different table types in a recent study of a full-service casual Italian restaurant:

<i>Table Location</i>	<i>Spending per Minute (SPM)</i>
Best tables (in the window)	\$.41
Worst tables (in the middle of the room)	\$.36
Overall average (all tables)	\$.40

Additional Sources

If you'd like to know more about dining room performance and analyses, you can access articles and assessment tools, many of them for free, at the Cornell University Center for Hospitality Research Web site at www.hotelschool.cornell.edu/chr/.

Here are some of the articles regarding measuring front-of-house performance:

- ✓ "Restaurant Revenue Management," by Sheryl E. Kimes, Ph.D., 2004 Cornell Hospitality Reports, Cornell University Center for Hospitality Research.
- ✓ "The Impact of Table Characteristics on Dining Duration and Spending," by Sheryl E. Kimes, Ph.D., and Stephani K.A. Robson, 2004, *Cornell Hotel and Restaurant Administration Quarterly*, Volume 45, No. 4, pages 333-34.
- ✓ "Dining Duration and Customer Satisfaction," by Breffni M. Noone, Ph.D., and Sheryl E. Kimes, Ph.D., 2005 Cornell Hospitality Reports, Cornell University Center for Hospitality Research.
- ✓ "Restaurant Capacity Effectiveness: Leaving Money on Tables," by Gary M. Thompson, Ph.D., 2007 Cornell Hospitality Reports, Cornell University Center for Hospitality Research.

So there's a difference of about 14 percent in spending per minute between the best and worst tables. While the dollar values in this particular analysis are pretty low, they can translate into quite a lot of money over time.

One possible explanation for this difference in spending and duration by table location lies in psychology: Most people prefer to position themselves next to architectural features as a way of establishing a personal territory and therefore controlling their privacy. You can see this at any fast-food place where people can choose to sit where they like: Most guests

will pick the tables around the perimeter of the room, or choose booths with low walls on at least two sides, whenever possible. Dining rooms that tend to make guests happier are ones where most if not all of the tables are somehow "anchored" by design. From your own restaurant's POS data, calculate your restaurant's SPMs by looking at spending and duration by table number. Is there a big difference in SPM between different kinds of tables? Maybe adjusting your table locations can help.

Of course, differences in server skill can also influence SPM, so doing a similar analysis that looks at SPM by server rather than table location can be a really helpful measure. If a server performs significantly below his or her peers, it's time to take a look at what might help. More training? Changing the section the server looks after? (You probably want to be rotating sections frequently, so that all servers get a chance at the "good" tables.) The SPM measure is an easy one to calculate but one that very few restaurateurs know about, despite its usefulness for the operator.

The Gang's All Here! (But Can You Seat Them?)

The final element in assessing dining room performance is working out whether your mix of tables makes sense for your clientele. Let's say that Friday and Saturday dinners at your restaurant are your busiest times, with all tables occupied for a good part of the night. This is good news ... as long as all the seats at each table are occupied. But if you have parties of two sitting at a four-top table, those two unsold seats are lost forever. And many restaurants underperform in this way: They have the wrong mix of tables for the party sizes that they serve.

During the design process, many operators make assumptions about the sizes of parties that will come to the restaurant on any given night and plan the dining room according to those assumptions, but very often those assumptions are incorrect. I recently

reviewed the floor plans for a proposed dining room in which all of the tables — with one lonely exception — were four-tops, and yet in the majority of restaurants of this particular type, the typical party size was just two people. So if the restaurant was built the way it was designed, the operator would have to seat parties of two at tables for four, resulting in a lot of lost revenue and in longer waits for parties of four.

So what mix of tables should you have for your dining room to perform at its best? A review of your POS data can once again assist you: The record for each check should tell you how many were in that particular party. (There is of course the challenge of split checks, which require a bit more detective work to unravel but are not impossible to identify. Multiple checks on the same date at the same table number that were closed within a few minutes of each other are your best clue to instances of split checks.) With this information, you can determine what proportion of your parties is single diners, twos, fours, and so on. Then compare these proportions with the mix of tables that your restaurant offers. Serious gaps are a good indicator that you should perhaps buy some new tables.

One approach to maximizing the potential of the dining room is to make all tables "two-tops," combining them as necessary to seat larger parties. While this approach is good in theory, as a practical matter, having to combine tables on the fly all the time puts more pressure on your staff. A better approach is indeed to use all two-top tables, but to leave some of them combined into fours most of the time, only separating them or combining them further when the need arises. How many to combine depends on the mix of party sizes that you see most frequently.

As an example, let's look at an 82-seat casual dining restaurant with a busy bar and a really great weekend dinner and Sunday brunch business that can result in a half-hour to 45-minute wait during peak times. The current

mix of tables includes 11 deuces and 15 four-tops. A review of this restaurant's POS data indicates that actual party sizes look more like this:

Party Size	Proportion of All Parties
1-2 people	70.2%
3-4 people	22.9%
5-6 people	4.5%
7 or more people	1.4%

So more than two-thirds of this restaurant's customers come in parties of two or less, and yet less than half of the tables are deuces. That means that a lot of the time, parties of two are being seated at a table that is too big for them, and the restaurant is giving up the possible revenue from those two extra seats at the table on busy nights when there is a wait for a table. Using the data above, let's see how much revenue that really is for a busy hour, assuming an average check of \$25:

Party Size	Proportion of all parties	Total # parties of this size on average	Parties of two at a two-top	Parties of two at a four-top	Seats at four-tops not used	Lost revenue (unused seats*av.check)
1-2 people	70.2%	26 tables * 70.2% = 18	11	18 - 11 = 7	7*2 = 14	14*\$25 = \$350

So by having too few deuces, this restaurant could be losing as much as \$350 in each turn of the dining room at busy times.

To do a similar calculation for your restaurant, take the party mix during your busiest meal period, because that's when having the right table mix will have the greatest effect on your revenues. Compare that party size mix with the mix of tables you have now: how many tables for two, four, six and so on. If you typically combine smaller tables to accommodate larger parties, just count the total number of tables that could be seated by separate parties: If a deuce and a four-top are typically laid out side by side as a six-top but *could* be separated into two tables, count them as one deuce and one four-top. Using the table above as a guide, figure out how many seats are going unused during your busy period because of a poor match between party size and table size.

Knowledge is Power

So to summarize, measuring your dining room performance — revenue per available seat hour, seat use, dining duration and table mix — gives you a wealth of information that helps you be a better operator. RevPASH calculations let you decide when to upsell and when not to. Seat use analysis tells you when you need to attract more customers, or when your demand is outstripping your supply of seats. Measuring dining duration helps you keep tabs on your service quality, while spending-per-minute calculations can help you fine-tune how your tables are arranged in the room. And, finally, comparing table mix with your party sizes helps you make sure you have the right combination of table sizes so that you aren't losing possible revenue.

But don't stop with just these measures. There's probably more that your POS system can tell you about your dining room performance than you see here. One restaurateur I know used his POS data to compare sales on sunny days versus cloudy ones, and found that his open-kitchen pizzeria did a lot better on cool cloudy days, prompting him to create a special "sunny day" promotion that got him plenty of free local press. Others use their POS data to evaluate server performance and set up friendly contests among staff members to drive higher average checks. Maybe you've established your own measures and tools for monitoring dining room performance — if you have, please drop me a note because I'd love to hear about it.

RS&G