Cornell Hospitality Quarterly
52(3) 241–252

DOI: 10.1177/1938965511398261
http://cqx.sagepub.com

Inaccuracy of the "Naïve Table Mix" Calculations

By Gary M. Thompson

Abstract

This article extends the analysis of table mixes for restaurants, on the principle that table mix helps drive revenue. The article presents the results of a simulation experiment of more than forty-six thousand restaurant contexts to evaluate the accuracy of "naïve table mix" calculations. An earlier study presented a simple method of calculating a recommended mix of tables in a restaurant, and another found that those calculations provided table mixes yielding within about 1 percent of the revenue provided by the optimal table mix. While the earlier simulation studies assumed that the space required by a table was directly proportional to the number of seats it contained, that is not always the case. Thus, this article presents space-oriented versions of the existing naïve table mix methods. The article also presents the results of a test of two other forms of the naïve table mix calculations, which include party revenue (or contribution). The simplest of the naïve table mix models proved to be the best, on balance. It yielded table mixes that generated more than 98 percent of the revenue achievable by near-optimal table mixes. As a caveat, however, in more than 5 percent of the simulated scenarios, the best naïve table mix model yielded less than 95 percent of the revenue achievable by a near-optimal table mix. These findings suggest that the naïve table mix calculations be used with caution.

Keywords

restaurants, revenue management, operations, simulation

A key aspect of revenue management is a facility's available capacity. In the restaurant context, an essential element of capacity is the restaurant's mix of tables. In a study with Sheryl Kimes, I demonstrated that a particular fast casual restaurant could increase its effective capacity by more than 30 percent simply by changing to a table mix that better matched its customer demand (Kimes and Thompson 2004). Generally speaking, matching capacity to demand means that the mix of table sizes in the restaurant corresponds to the restaurant's mix of sizes of customer parties. While a 30 percent capacity increase may be on the high side for most restaurants, in a subsequent study of sixty-eight restaurants, I found an average increase of more than 14 percent in effective capacity with an improved table mix. In sum, optimizing a restaurants' mix of tables seems highly worthwhile.

The question that has arisen in the course of these table mix studies regards the quality of mixes identified by simple mathematical calculations. In a 2002 study, I presented the simplest approach for identifying table mixes, which I called the "naïve table mix" (Thompson 2002). This calculation uses information on the probability of different size parties arriving at the restaurant, the size of tables at which the parties are seated, and the overall seating capacity of the restaurant. However, this model did not account for the facts that mean dining durations are not identical across party sizes and the space required by a table is not always proportional to its number of seats. In another study with Sheryl Kimes, however, we found that my original method yielded revenue that exceeded 99 percent of the revenue achieved with a near-optimal table mix. Not only that, but the naïve table mix calculation performed better than a duration-modified method we tested (Kimes and Thompson 2005). These results suggest that managers would be well served using the simple models.

In this article, I discuss a study that goes further than that, however, to evaluate the effectiveness of two new forms of the naïve table mix calculations, which include information on the revenue (or contribution) of each party size. In this study, I evaluated the effectiveness of the four forms of the naïve table mix in more than forty-six thousand simulated restaurant contexts, differing on thirteen factors.

A puzzlement. The four naïve table mix approaches returned mixed results. Overall, the best of the naïve methods yields table mixes having within 2 percent of the revenue of near-optimal table mixes, which is poorer performance than has been previously observed. Indeed, in more than 5 percent of the simulated scenarios, the best naïve table mix yielded less than 95 percent of the revenue achievable by a near-optimal table mix. Most troubling, though, was that in a tiny fraction of the scenarios the revenue of the best naïve table mix was off by more than 10 percent, and these cases of poor performance were clustered under a certain set of restaurant conditions. So, contrary to my earlier conclusions, these findings suggest that restaurateurs should employ sophisticated methods for identifying table mixes if they can. Having said that, I also demonstrate the value of the simplest of table-mix calculations.

Review of Relevant Literature

Having developed the concept of restaurant revenue management (RRM) in 1998, Kimes and her coauthors also defined RevPASH, or revenue per available seat hour, as a metric of restaurant performance that captures the effects of revenue, time, and capacity (measured in seats) (Kimes et al. 1998). Research on RRM can be categorized based on whether it focuses on capacity or demand. To be thorough, this review includes some of the studies I just mentioned.

In an earlier simulation, I examined whether restaurants perform better when they have a mix of table sizes matched to specific party sizes, or when they have tables that can combined to seat larger parties (Thompson 2002). Then I considered how combinable tables can be situated to yield the best performance (Thompson 2003) and, with Sheryl Kimes, I evaluated methods of determining a restaurant's table mix, using both a simulation study and data from a real restaurant (Kimes and Thompson 2005). Applying this approach to a fast casual restaurant, we found that it could increase its effective capacity by more than 30 percent by matching its table mix with customer demand (Kimes and Thompson 2004). In a study of sixty-eight restaurants, I found that changing the table mix would increase capacity by more than 14 percent, on average (Thompson 2007).

As I stated above, my 2002 study first presented the naïve table mix model, and then Kimes and I modified that model to incorporate dining duration by party size. In that study we assessed two simulated annealing search heuristics that each evaluated one hundred different table mixes during the search process. We took the better performing of the two and, using data from a real restaurant, we found table mixes that yielded revenue within 0.02 percent of that yielded by the optimal table mixes. Because this particular heuristic, called SimAnnealSearch, yielded such strong performance in a fraction of the time required to evaluate all possible table mixes, we could use it as a benchmark for the simpler approaches in a simulation of ninety-six restaurant contexts. For those ninety-six contexts, we found that the initial naïve-seats model yielded revenue that was within 1 percent of that given by SimAnnealSearch.

A number of simulation-based studies have been conducted related to restaurant operations to analyze, among other things, table size, table mix, and combinability of tables (e.g., Field, McKnew, and Kiessler 1997; Thompson 2002; Thompson 2003; Kimes and Thompson 2005; Thompson and Kwortnik 2008; Thompson 2009; and Thompson and Sohn 2009).

Naïve Table Mix Models

I will start with my 2002 naïve-seats model, followed by the modification that Kimes and I made in 2005. Then I will develop the space-based modifications of these models. For ease of presentation, I will first define all parameters used in the models:

b_p = probability of party of size p;
$Duration_p$ = mean dining duration for parties of size p;
$Value_p$ = mean value (revenue or contribution) for parties of size p;
$PartyFits_i$ = party sizes that will fit in and be assigned to table size i;
$Space_i$ = space required by table size i;
$Tables$ = set of allowable table sizes;
$TotalSeats$ = total seats in the restaurant; and
$TotalSpace$ = total space available in the restaurant for seating.

The original version of the naïve table mix model (Naïve-Seats; Thompson 2002) specified the ideal number of tables in the restaurant as

$$IdealTables_i = \left(TotalSeats * \frac{WeightedSeats_i}{\sum_{j \in Tables} WeightedSeats_j} \right) \div i, \qquad (1)$$

where

$$WeightedSeats_i = i * \sum_{p \in PartyFits_i} b_p . \qquad (2)$$

Notice that the only parameters in naïve-seats are the party size probabilities, the number of seats in each allowable table size, the matching of parties to tables, and the total number of seats available in the restaurant.

The modified version of the original table mix model (Naïve-SeatsDuration; Kimes and Thompson 2005) was designed to incorporate differences in dining duration across party sizes:

$$IdealTables_i = \left(TotalSeats * \frac{WeightedSeatMinutes_i}{\sum_{j \in Tables} WeightedSeatMinutes_j} \right) \div i, \qquad (3)$$

where

$$WeightedSeatMinutes_i = i * \sum_{p \in PartyFits_i} \left(b_p * Duration_p \right). \qquad (4)$$

Neither Naïve-Seats nor Naïve-SeatsDuration considers that tables may require space that is not proportional to the number of seats in the table. Thus, the first step for this study is to relax this potentially limiting assumption, which requires that the item being allocated is total space, not total seats. The space-oriented version of Naïve-Seats, *Naïve-Space*, is

$$IdealTables_i = \left(TotalSpace * \frac{WeightedSpace_i}{\sum_{j \in Tables} WeightedSpace_j} \right) \div Space_i , \qquad (5)$$

where

$$WeightedSpace_i = Space_i * \sum_{p \in PartyFits_i} b_p . \qquad (6)$$

Likewise, the space-based modification of Naïve-SeatsDuration, *Naïve-SpaceDuration*, is

$$IdealTables_i = \left(TotalSpace * \frac{WeightedSpaceMinutes_i}{\sum_{j \in Tables} WeightedSpaceMinutes_j} \right) \div Space_i , \qquad (7)$$

where

$$WeightedSpaceMinutes_i = Space_i * \sum_{p \in PartyFits_i} \left(b_p * Duration_p \right). \qquad (8)$$

Next, I will incorporate information on the revenue (or contribution) yielded by each size of party, since the amount spent per person typically declines as party size increases (see, for example, Kimes and Robson 2004; Kimes and Thompson 2004). Because ignoring revenue in the table mix calculations could lead to inferior table mixes, this model incorporates revenue (or contribution) into *Naïve-SpaceDuration*, yielding *Naïve-Space DurationValue*:

$$IdealTables_i = \left(TotalSpace * \frac{WeightedSpaceValueMinutes_i}{\sum_{j \in Tables} WeightedSpaceValueMinutes_j} \right) \div Space_i , \qquad (9)$$

where

$$WeightedSpaceValueMinutes_i = Space_i * \sum_{p \in PartyFits_i} \left(b_p * Duration_p * Value_p \right). \qquad (10)$$

A different way of considering revenue (or contribution) is to base the space allocated to each table size using the relative revenue (or contribution) that arises per space-minute for parties seated at that size table. This model, *Naïve-ValuePerSpaceMinute*, is

$$IdealTables_i = \left(TotalSpace * \frac{WeightedValuePerSpaceMinute_i}{\sum_{j \in Tables} WeightedValuePerSpaceMinute_j} \right) \div Space_i , \qquad (11)$$

where

$$WeightedValuePerSpaceMinute_i = \sum_{p \in PartyFits_i} \left(b_p * \frac{Value_p}{Space_i * Duration_p} \right). \qquad (12)$$

An issue with any of the naïve methods is dealing with fractional numbers of tables. Kimes and I used an integer programming model in the 2005 study to round fractional tables to yield the least total discrepancy across all table sizes between unrounded and rounded numbers of seats. In this article, I use a simple rounding heuristic: I start with the largest size table and move sequentially to the smallest size table, to reflect the fact that flexibility increases as table size decreases. So, I round the largest size table from its ideal number to the nearest whole number. I then repetitively move to the next largest size table and round it. However, I accumulate any discrepancies in space utilization as the analysis moves to the smaller tables. For example, if I round up a particular table size, I have a net overuse of space, which I can then use to adjust downward the ideal number of the next-smaller-size table. The advantages of this sequential approach are that it is easy, fast, and does not require integer-programming software. I believe that this sequential rounding approach better mimics how restaurateurs would apply the naïve methods.

Design of the Simulation Experiment

To test the various forms of the naïve table mix calculations under a wide variety of restaurant scenarios, this simulation experiment contained the thirteen factors summarized in Exhibit 1. Nine factors were environmental, or outside the control of restaurant managers. Of the remaining four, two were design-related and two were related to operational decisions. By including the design- and operational-related factors, I have the opportunity to develop prescriptive guidelines for the restaurant industry, by comparing the revenue achieved across the levels of the design- and operational-related factors.

Environmental factors. The nine environmental factors were restaurant size; mean party size; average check variation across and within party sizes; customers' propensity to wait for a table; duration variation across and within party sizes; demand intensity; and length of the peak demand period. Restaurant size had three levels, corresponding to the ability to seat approximately 10, 30, and 100 parties simultaneously (corresponding to having approximately 10, 30, and 100 tables, which I categorized as small, medium, and large). It is preferable to define size based on number of parties, rather than number of seats, because we wish to avoid confounding effects with the table space requirements experimental factor. These three size levels match

Exhibit 1:
Experimental Factors in the Simulation Experimental, Categorized by Type

	Abbreviation	Number of Levels: Levels
Environmental factor		
Restaurant size	RstSz	3: Approximately 10, 30, and 100 parties seated simultaneously
Mean party size	MnPtySz	2: 2.5 and 3.0 people
Average check variation across party sizes	ACVaPS	2: $10 average spent per person for parties of 1, decreasing by $0.10 and $0.20 per additional person in the party
Average check variation within party sizes	ACVwPS	2: Coefficients of variation of 0.15 and 0.30
Propensity to wait	PrToWt	2: Short and long: wait tolerances of 15-30 minutes and 30-60 minutes for parties of 1-10, respectively
Duration variation across party sizes	DVaPS	2: Ratio of duration for a party of 10 compared to the duration for a party of 1 equal to 1.5 and 2.0
Duration variation within party sizes	DVwPS	2: Coefficients of variation of 0.20 and 0.35
Demand intensity	DmdInt	3: 100%, 115%, and 130% of restaurant capacity
Peak demand length	PkDmdLen	5: 2, 2.5, 3, 3.5, 4, 4.5, and 5 hours
Design factor		
Table space proportions	TblSpPrp	2: Space requirements proportion and nonproportional to the number of seats
Table size options	TblSzOpts	2: Only tables with even-numbered seats; tables with both odd- and even-numbered seats
Operational factor		
Table assignment rule	TblAssRl	2: Largest party that fits; longest waiting party that fits
Right-sizing rule	RightSzRl	2: Tighter and looser right-sizing: maximum table sizes of (2, 2, 4, 4, 6, 6, 8, 8, 10, 10) and (4, 4, 6, 6, 8, 8, 10, 10, 10, 10) for party sizes of 1 to 10, respectively

well with size ranges used in earlier studies (an overall range of 50 to 450 seats). We did not include exceptionally large restaurants, for the simple reason that simulation run-time is greatly affected by a restaurant's size. Including a restaurant size of 300 parties (tables) would have tripled the time required to run the experiment, which required more than three weeks of computer time with the existing parameters.

Mean party size had two levels: 2.5, people which is about the reported industry average (Kimes and Robson 2004); and 3.0 people, which I included for thoroughness (see Exhibit 2).

Average check variation across party sizes had two levels (also in Exhibit 2), both of which were selected to reflect the common occurrence that larger parties tend to spend less per person than smaller parties. The differences in average check across party sizes ($0.10 and $0.20 for a one-person difference in party size) were the same as those I used in a study with Heeju Sohn (2009), this time using a base check value of $10 per person. I used similar values for average check variation within party sizes (two levels) as that 2009 study.

The two levels for customers' propensity to wait for a table were based on anecdotal evidence suggesting that larger parties are willing to wait longer for a table than are smaller parties (see Exhibit 2). Individual parties were assigned a willingness to wait, based on their party size, using a normally distributed random variate, where the

Exhibit 2:
Parameter Levels, by Level of the Experimental Factors Related to Party Size

	Party Size Probabilities, by Mean Party Size		Average Check per Person, by Level of Variation		Average Wait Tolerance, in Minutes, by Level of Tolerance	
Party Size	2.5	3.0	Low	High	Low	High
1	0.2	0.12	$10.00	$10.00	15.00	30.00
2	0.47	0.32	$9.90	$9.80	16.67	33.33
3	0.15	0.26	$9.80	$9.60	18.33	36.67
4	0.1	0.18	$9.70	$9.40	20.00	40.00
5	0.03	0.05	$9.60	$9.20	21.67	43.33
6	0.02	0.03	$9.50	$9.00	23.33	46.67
7	0.01	0.02	$9.40	$8.80	25.00	50.00
8	0.01	0.01	$9.30	$8.60	26.67	53.33
9	0.005	0.005	$9.20	$8.40	28.33	56.67
10	0.005	0.005	$9.10	$8.20	30.00	60.00

Exhibit 3:
Space Required per Table, by Number of Table Seats

Metric	Naïve-Space	Naïve-SpaceDuration	Naïve-SpaceDurationValue	Naïve-ValuePerSpaceMinute
Average percentage revenue shortfall	1.86	2.81	11.72	10.88
Percentage of instances yielding the best naïve table mix	52.54	49.52	2.09	0.86

coefficient of variation was 0.20. In earlier research (Thompson 2002), I used 90 minutes for all parties and three levels for willingness to wait, equal to 1/3, 2/3, and 1 times the mean dining duration by party size.

Duration variation across party sizes had two levels, based on the observation that dining duration increases with an increase in party size (Bell and Pliner 2003; Kimes and Robson 2004; Kimes and Thompson 2005). The two levels of this factor set the ratio of the dining time for parties of 10 people to 1.5 and 2.0 times that of the dining duration for parties of 1 person, consistent with a recent study of mine (Thompson 2009). For both factor levels, there was a straight linear relationship between mean dining duration and party size.

Duration variation within party sizes had two levels: coefficients of variation in duration, by party size, of 0.20 and 0.35. These values fall within the range of 0.16 to 0.50 reported in the literature (Bell and Pliner 2003; Kimes and Robson 2004).

I applied a different approach to demand intensity than in the past. This time I applied three levels, starting with a "low" level in which demand was sufficient to result in the restaurant running at full capacity during the peak period. "Medium" inflated customer demand by 15 percent, and "high" was 30 percent higher than the base. As such, this factor mimics situations that exist in popular restaurants, where demand exceeds capacity. In other simulations, I have used assumptions ranging from 95 percent of capacity (Thompson 2003; Thompson and Sohn 2009) to a high of 120 percent of capacity (Kimes and Thompson 2005).

The length of the peak demand period had five levels: 2, 2.5, 3, 3.5, and 4 hours, reflecting a broad range of restaurant settings. This range falls within the range of peak durations that have been used in the literature (e.g., Field, McKnew, and Kiessler 1997: 5 and 7 hours; Thompson 2009: 9 peak durations, ranging from 1 to 5 hours in 30-minute increments).

Design factors. The experiment included two factors related to restaurant design issues: table space requirements and table size options. For table space requirements, I applied one level where each table required space proportional to its number of seats and a second level in which the space required decreased on a per-seat basis (see Exhibit 3). This is a departure from my previous table mix studies, which assumed that the space requirements are directly proportional to the number of seats (Thompson 2002, 2003, 2009; Kimes and Thompson 2005). The second factor related to restaurant design is table size. The first level of this factor allows only tables with even-numbered seats, as has occurred in all previous studies (Field, McKnew, and Kiessler 1997; Thompson 2002, 2003, 2009; Kimes and Thompson 2005), while the second level allowed any size table between one and ten seats.

Operational factors. I included two operational factors in the simulation experiment: a table assignment rule and a maximum table size for party assignments. The table assignment rule had two levels: assigning an available table to the largest waiting party that fits in the table regardless of wait (largest) and assigning an available table to the party that has been waiting the longest and that fits in the table (longest). Previous research has used the largest rule (Thompson 2002, 2003, 2009; Kimes and Thompson 2005), with one exception (Field, McKnew, and Kiessler 1997), where the longest rule was used.

Right-sizing, a term that we believe was first used by Kimes (2004), has two levels: tighter right-sizing, which allows a party to be assigned to a table no bigger than the smallest even-sized table in which the party fits; and looser right-sizing, which allows a party to be assigned to a table no bigger than the smallest even-sized table in which the party fits plus two seats. My previous studies did not take this consideration into account.

Simulation assumptions. None of the assumptions applied here are nonstandard, since similar assumptions have been made in previous studies. In fact, because of the extensive scope of our simulation study, this study actually has fewer assumptions than have been made in earlier studies. We made the following eight assumptions:

1. Parties ranged from one to ten people.
2. Parties would not be split across tables.
3. Tables could not be combined.
4. All demand was from walk-ins.
5. Arrivals were based on a Poisson distribution, which mimics the typical arrival process in restaurants in that several parties often arrive close

together, followed by notable time gaps without party arrivals (resulting in a negative exponential distribution of the time between party arrivals).

6. Dining duration had a log normal distribution, mimicking the situation of a few parties lingering over their meals.
7. The number of parties that could be waiting had no limit, so the only determinant of whether a party would be served was whether a table became available within the party's wait tolerance.
8. Customer demand ramped up for an hour prior to the peak period, followed by an hour-long ramp-down after the peak.

Simulation process. I applied a full factorial experimental design for the thirteen factors and their levels, resulting in a total of 46,080 simulated restaurant contexts. That is, 3 levels of restaurant size × 2 levels of restaurant size × 2 levels of mean party size × 2 levels of average check variation across party sizes × 2 levels of average check variation within party sizes × 2 levels for customers' propensity to wait × 2 levels of duration variation across party sizes × 2 levels of duration variation within party sizes × 3 levels of demand intensity × 5 levels of peak demand length × 2 levels of table space proportions × 2 levels of table size options × 2 levels of table assignment rules × 2 levels of right-sizing rules. For each of the 46,080 combinations of the factor levels, I simulated and stored all relevant information for one hundred days of operation of the restaurant. This information included, for each party, its size, arrival time, dining duration, average check, and willingness to wait. Storing the information for a scenario allowed me to control for extraneous variation in the results, because each method was evaluated under common conditions. Again, the models are Naïve-Space, Naïve-SpaceDuration, Naïve-SpaceDurationValue, and Naïve-ValuePerSpaceMinute. I identified each model's recommended table mixes. I then simulated the restaurant using each table mix, using the common, stored data, and ran a heuristic, much like the SimAnnealSearch heuristic, to search for the best possible mix of tables for that particular restaurant context. I used the SimAnnealSearch because it came exceedingly close to optimal (I will refer to it as "near-optimal"). Given the size of this experiment, an exhaustive search for the best table mix would have been impractical. For speed, I coded the simulation model itself in Microsoft's Visual Basic 6 and compiled it as an executable program. I ran the experiment on several computers simultaneously, but on a single computer it would have taken 24 days (573.4 hours, assuming a Xeon 2.8 Ghz processor).

Results of the Simulation Experiment

To evaluate the naïve mix methods, I calculated a percentage revenue shortfall, which represents the revenue shortfall arising from a table mix given by a naïve method, as compared to the near-optimal table mix. Exhibit 4 reports the overall results across the 46,080 simulated restaurant contexts. In terms of the average percentage revenue shortfall, the simplest estimate, Naïve-Space, proved to be the most accurate, with a revenue shortfall of just under 2 percent. These findings echo those I reported with Sheryl Kimes (2005), but this time the revenue shortfall percentage was easily double that in the 2005 study. This finding is perhaps related to the much broader set of restaurant environments simulated in this study. The two approaches incorporating revenue—Naïve-SpaceDurationValue and Naïve-ValuePerSpaceMinute—both performed poorly, having revenue shortfalls of more than 10 percent on average. As for Naïve-SpaceDuration, its average percentage revenue shortfall was more than 50 percent higher than that of Naïve-Space. However, those two methods were quite similar in the frequency with which they yielded the best of the Naïve table mixes, a situation that occurred for both in approximately half of the settings evaluated.

Exhibit 5 presents the percentage of scenarios in which each of the four methods yielded revenue within a specified tolerance of the near-optimal revenue, starting at 99 percent, and declining to 90 percent or less. The table shows that Naïve-Space was superior to the other three methods, and it yielded revenue within 2 percent of near-optimal in close to two-thirds of the scenarios; and the percentage of scenarios where it was in the 90-percent or less category was small, just 0.4 percent of the scenarios. Naïve-SpaceDuration was reasonably similar in performance to Naïve-Space, but both Naïve-SpaceDurationValue and Naïve-ValuePerSpaceMinute were notably poorer in performance.

Exhibit 6 presents the percentage revenue shortfalls of the four methods, by level of the experimental factors. Examining Naïve-Space, one sees that it performed more poorly in the smallest restaurant size, the lowest-demand intensity, when table sizes were nonproportional, and when all size table options were possible. For the other factors, differences in the performance of Naïve-Space were less pronounced.

Exhibit 7 reports the frequency with which each method yielded the best table mixes, by level of the experimental factors. Naïve-Space yielded the greatest percentage of best table mixes under the small and medium restaurant sizes, larger mean party sizes, both levels of average check variation across and within party sizes, short propensity to wait, larger variation in dining durations across party sizes, both levels of dining duration variation within party sizes, medium and high demand intensity, peak demand durations of more than two hours, proportional table sizes, both levels of table size options, assigning parties to tables based on the longest wait, and with tighter right-sizing. Naïve-SpaceDuration yielded the greatest percentage of best table mixes under the large restaurant size, smaller mean party sizes, longer

Exhibit 4:
Overall Results

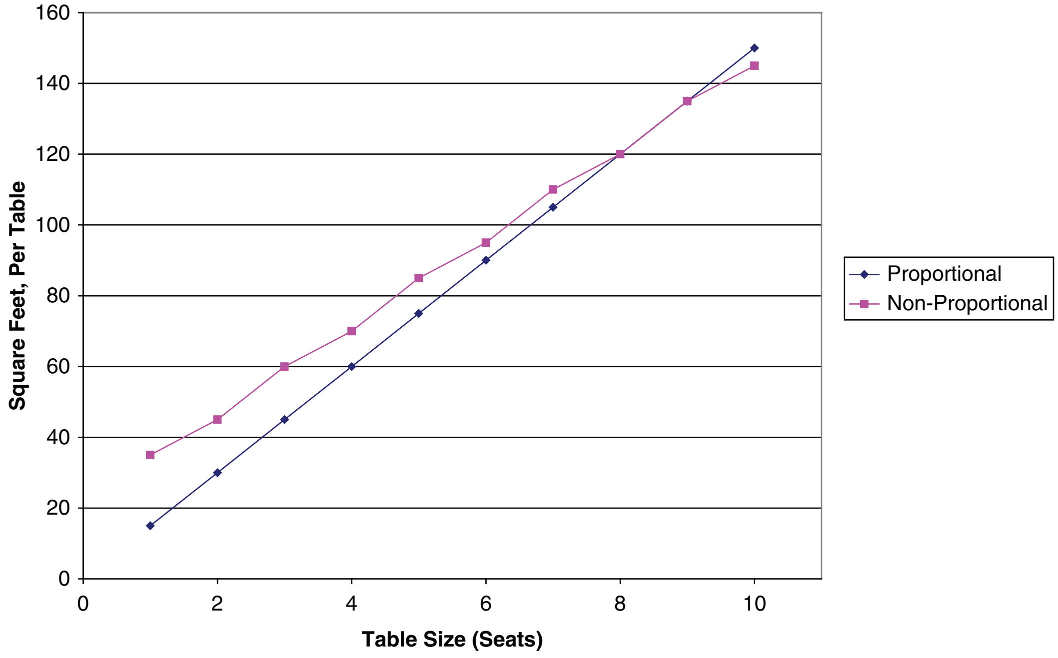

Exhibit 5:
Percentage of Scenarios in Which the Naïve Methods Yielded Revenue within a Specified Tolerance of the Near-Optimal Revenue

Percentage of Near-Optimal Revenue	Naïve-Space	Naïve-SpaceDuration	Naïve-SpaceDurationValue	Naïve-ValuePerSpaceMinute
99 or more	12.50	12.63	0.25	0.50
98 or more	64.08	58.04	2.01	2.67
97 or more	79.77	69.55	4.43	5.18
96 or more	89.61	76.41	7.78	8.48
95 or more	94.67	80.34	11.97	12.22
90 or less	0.43	5.87	61.02	52.33

propensity to wait, smaller variation in dining durations across party sizes, low demand intensity, peak demand durations of two hours, nonproportional table sizes, assigning parties to tables based on the largest size, and with looser right-sizing. The other two performed relatively poorly by comparison.

Exhibit 8 reports the results of regression models for each method, using percentage revenue shortfall as the dependent variable. The results for Naïve-Space show that its revenue shortfall was statistically significant, and higher, with smaller restaurant sizes, smaller mean party sizes, longer propensity to wait, larger variation in dining duration across party sizes, lower demand intensity, shorter durations of the peak demand window, nonproportional table sizes, all table-size options, assigning parties to tables based on the "largest" rule, and with tighter right-sizing.

Exhibit 9 reports information on the worst-case performance of Naïve-Space—the 196 instances (just 0.4 percent of the total) where its revenue shortfall exceeded 10 percent. There is definite clustering of the poor performance scenarios by certain levels of the experimental factors. Let us look, for example, at the 480 test scenarios that involved these characteristics: smallest restaurant size, larger mean party size, smaller variation in dining duration across party sizes, proportional table sizes, odd and even table-size options, and tighter right-sizing. Naïve-Space returned a revenue shortfall of more than 10 percent in 187 (39 percent) of those 480 scenarios.

Exhibit 6:
Percentage Revenue Shortfall, by Level of the Experimental Factors

Factor	Level	Naïve-Space	Naïve-SpaceDuration	Naïve-SpaceDurationValue	Naïve-ValuePerSpaceMinute
RstSz	10	2.42	5.38	12.52	8.41
	30	1.59	2.21	11.22	11.12
	100	1.58	0.84	11.41	13.09
MnPtySz	2.5	2.03	2.68	13.93	10.94
	3.5	1.70	2.94	9.50	10.81
ACVaPS	Smaller	1.88	2.78	11.74	10.83
	Larger	1.85	2.84	11.69	10.92
ACVwPS	Smaller	1.86	2.81	11.71	10.87
	Larger	1.87	2.82	11.72	10.88
PrToWt	Shorter	1.83	2.96	11.72	10.17
	Longer	1.90	2.67	11.71	11.58
DVaPS	Smaller	1.76	2.36	10.71	9.90
	Larger	1.96	3.26	12.72	11.85
DVwPS	Smaller	1.87	2.82	11.75	10.85
	Larger	1.86	2.80	11.68	10.90
DmdInt	100	2.17	2.37	9.79	11.93
	115	1.68	2.61	11.96	11.02
	130	1.75	3.45	13.40	9.67
PkDmdLen	2 hours	1.92	2.66	11.18	11.16
	2.5 hours	1.88	2.75	11.53	10.97
	3 hours	1.85	2.82	11.79	10.85
	3.5 hours	1.84	2.89	11.97	10.74
	4 hours	1.83	2.94	12.12	10.66
TblSpPrp	Proportional	1.34	2.96	14.38	14.16
	Nonproportional	2.38	2.66	9.05	7.59
TblSzOpts	Even only	1.20	1.97	12.05	9.69
	Odd and even	2.52	3.65	11.38	12.06
TblAssRl	Largest	1.90	2.65	10.76	11.09
	Longest	1.82	2.98	12.67	10.66
RightSzRl	Tighter	1.93	3.83	14.35	11.38
	Looser	1.80	1.79	9.09	10.37

Exhibit 10 presents information on the relative revenue of the near-optimal table mixes, for the design- and operational-related experimental factors. This table indicates that higher overall revenue was attained with proportional table space, both odd and even table sizes, seating parties based on the "largest" rule, and with tighter right-sizing. The relative revenue increase was the greatest for the right-sizing operational rule, where the tighter application of right-sizing increased revenue by an average of 4.48 percent compared to the looser application of right-sizing. The second largest difference in revenue arose with the table size options, where the use of table with both odd- and even-number seats yielded 2.43 percent more revenue than using only tables with even-numbered seats.

Discussion

An interesting question is why the simplest of the naïve methods—the one that uses only party probabilities, table sizes, and the total space of the restaurant—outperformed the three other naïve methods, which all incorporate additional—and seemingly relevant—information. I believe the answer lies in which customers are served, and which are lost, as a result of the table mixes identified by each method. An optimal table mix will allocate the space in the restaurant in a way that maximizes revenue, considering all the environmental, design-related, and operational-related factors. This will mean that greater numbers of the less valuable customers will be turned away. Examining which

Exhibit 7:
Percentage of Instances Where Each Naïve Method Yielded the Best Table Mix, by Level of the Experimental Factors

Factor	Level	Naïve-Space	Naïve-SpaceDuration	Naïve-SpaceDurationValue	Naïve-ValuePerSpaceMinute
RstSz	10	75.01	32.27	5.13	2.59
	30	57.47	41.65	0.89	0.00
	100	25.13	74.64	0.24	0.00
MnPtySz	2.5	48.67	58.99	0.78	0.81
	3.5	56.41	40.04	3.39	0.92
ACVaPS	Smaller	51.76	50.16	2.58	0.63
	Larger	53.32	48.87	1.59	1.10
ACVwPS	Smaller	52.63	49.57	2.05	0.82
	Larger	52.44	49.47	2.12	0.91
PrToWt	Shorter	56.85	45.32	1.58	1.18
	Longer	48.22	53.72	2.60	0.55
DVaPS	Smaller	48.00	56.61	2.16	1.36
	Larger	57.07	42.42	2.01	0.36
DVwPS	Smaller	52.94	49.11	2.03	0.88
	Larger	52.13	49.92	2.14	0.85
DmdInt	100	36.48	63.83	4.73	0.06
	115	50.51	52.98	1.19	0.66
	130	70.63	31.73	0.34	1.87
PkDmdLen	2 hours	47.48	53.74	3.29	0.55
	2.5 hours	50.90	50.95	2.45	0.75
	3 hours	53.27	49.10	1.86	0.85
	3.5 hours	54.89	47.54	1.52	1.00
	4 hours	56.14	46.25	1.31	1.17
TblSpPrp	Proportional	67.98	35.82	0.13	0.15
	Nonproportional	37.10	63.21	4.04	1.58
TblSzOpts	Even only	53.41	51.36	1.42	1.71
	Odd and even	51.66	47.67	2.75	0.02
TblAssRl	Largest	47.19	54.54	2.94	0.47
	Longest	57.88	44.49	1.23	1.25
RightSzRl	Tighter	61.50	44.28	0.01	0.23
	Looser	43.58	54.75	4.16	1.49

customers are lost with each method, compared to the numbers lost in the near-optimal table mixes, should provide insight into the net effectiveness of the different methods. Exhibit 11 provides such a comparison. It shows that, across the entire experiment, Naïve-Space comes closest to the near-optimal table mixes in allocating space, since the customers lost is the closest match. Naïve-SpaceDuration allocates more space to larger parties than is ideal, since fewer customers are lost from larger parties, but this comes at the expense of a higher loss of customers from smaller parties. Naïve-SpaceDurationValue allocates far too much capacity to large parties, which results in a large loss of customers from smaller parties. Interestingly, Naïve-ValuePerSpaceDuration allocates too much capacity to the parties of one and two but consequently underserves all other party sizes.

Like any study, this one has limitations. The study's eight assumptions map well onto real restaurant situations, but because these simulations used a large number of experimental factors, most of which had only two factor levels, I urge caution in blindly applying the findings. For restaurateurs, the safest course of action would be to conduct an experimental study based on their specific restaurant setting. In applying a critical eye to the study, with insight gained from its findings, my largest concern relates to party size distributions. This simulation used just two levels, but there are an infinite number of party size distributions. As a research extension, then, I suggest repeating the experiment,

Exhibit 8:
Regression Results for Predicting Percentage Revenue Shortfall, for Each Naïve Method

Factor	Naïve-Space	Naïve-SpaceDuration	Naïve-SpaceDurationValue	Naïve-ValuePerSpaceMinute
Intercept	3.0649***	−3.1893***	8.0340***	12.1574***
RstSz	−0.0070***	−0.0424***	−0.0084***	0.0458***
MnPtySz	−0.3338***	0.2526***	−4.4315***	−0.1248***
ACVaPS	−0.2643	0.6779**	−0.5206*	0.8418***
ACVwPS	0.0911	0.0550	0.0242	0.0670
PrToWt	0.0713***	−0.2917***	−0.0050	1.4117***
DVaPS	0.4044***	1.8089***	4.0103***	3.9067***
DVwPS	−0.0265	−0.1687	−0.4196**	0.2986
DmdInt	−1.3970***	3.6250***	12.0291***	−7.5282***
PkDmdLen	−0.0421***	0.1404***	0.4615***	−0.2470***
TblSpPrp	1.0383***	−0.3049***	−5.3278***	−6.5697***
TblSzOpts	1.3174***	1.6829***	−0.6716***	2.3627***
TblAssRl	−0.0809***	0.3301***	1.9151***	−0.4236***
RightSzRl	−0.1350***	−2.0398***	−5.2620***	−1.0166***
R^2	.2596	.4363	.7893	.7353
Adjusted R^2	.2594	.4361	.7893	.7352
Standard error	1.5586	2.5222	2.4965	2.5424

*Significant at p = .05 level. **Significant at p = .01 level. ***Significant at p = .001 level.

Exhibit 9:
Occurrence Frequency of More Than 10 Percent Revenue Shortfall of Naïve-Space, by Level of the Experimental Factors

Factor	Level	Number of Occurrences	Factor	Level	Number of Occurrences
RstSz	10	196	DmdInt	100%	70
	30	0		115%	66
	100	0		130%	60
MnPtySz	2.5	0	PkDmdLen	2 hours	41
	3.5	196		2.5 hours	41
ACVaPS	Smaller	94		3 hours	42
	Larger	102		3.5 hours	40
ACVwPS	Smaller	101		4 hours	32
	Larger	95	TblSpPrp	Proportional	196
PrToWt	Shorter	86		Nonproportional	0
	Longer	110	TblSzOpts	Even	0
DVaPS	Smaller	187		Odd and even	196
	Larger	9	TblAssRl	Largest	89
DVwPS	Smaller	92		Longest	107
	Larger	104	RightSzRl	Tighter	196
				Looser	0

Exhibit 10:
Relative Revenue of the Design- and Operational-Related Experimental Factors

Factor	Ratio of Levels	Average Revenue Increase
TblSpPrp	Proportional/nonproportional	0.34%
TblSzOpts	Odd and even/even only	2.43%
TblAssRl	Largest/longest	1.21%
RightSzRl	Tighter/looser	4.48%

locking in a level for any factors to which performance was insensitive and expanding the range of levels for those factors to which the performance of the naïve methods were sensitive.

It goes without saying that regardless of whether a naïve method or a sophisticated simulation model is used, accurate data are a necessity. As I pointed out above, Naïve-Space, the most accurate of the four naïve methods tested, requires

Exhibit 11:
Daily Customers Lost, by Party Size, by Naïve Method

Party Size	Naïve-Space	Naïve-SpaceDuration	Naïve-SpaceDurationValue	Naïve-ValuePerSpaceDuration	Near Optimal
1	2.71 (3.6%)	4.97 (6.7%)	18.55 (15.9%)	0.30 (0.2%)	4.15 (5.9%)
2	4.62 (6.1%)	8.98 (12.0%)	54.29 (46.6%)	0.66 (0.5%)	5.23 (7.5%)
3	13.35 (17.5%)	13.85 (18.5%)	20.17 (17.3%)	23.17 (18.2%)	11.20 (16.0%)
4	10.13 (13.3%)	10.33 (13.8%)	11.61 (10.0%)	21.76 (17.1%)	7.24 (10.4%)
5	10.02 (13.2%)	8.42 (11.3%)	3.15 (2.7%)	20.59 (16.2%)	8.56 (12.2%)
6	7.65 (10.0%)	6.95 (9.3%)	2.60 (2.2%)	16.07 (12.7%)	5.88 (8.4%)
7	8.45 (11.1%)	5.42 (7.3%)	1.36 (1.2%)	14.97 (11.8%)	8.31 (11.9%)
8	6.99 (9.2%)	5.28 (7.1%)	1.21 (1.0%)	11.57 (9.1%)	6.51 (9.3%)
9	6.04 (7.9%)	4.80 (6.4%)	1.60 (1.4%)	8.49 (6.7%)	6.16 (8.8%)
10	6.22 (8.2%)	5.71 (7.6%)	1.96 (1.7%)	9.42 (7.4%)	6.73 (9.6%)
Total	76.17	74.72	116.51	127.00	69.97

Note: Numbers are the average number of customers lost per day across all 46,080 simulated environments. Numbers in parentheses are number of customers lost for a party size, expressed as a percentage of all customers lost by that method.

the least data: only party size probabilities, table space requirements, and the space available in the restaurant. It is important to note that the party size probabilities should be based on unconstrained demand, rather than constrained demand. Constrained demand—what is recorded in the restaurant's point of sale system—is affected by the restaurant's existing mix of tables, and so it can be distorted and consequently should not be used on its own. By contrast, unconstrained demand is the true mix of party sizes for a restaurant (see Orkin 1998). The best way to collect this information is recording information on the customers that were lost and then integrate those numbers with data on customers served.

Conclusions

Overall, the performance of the naïve methods must be judged as quite good, given that they are all simple, back-of-the-envelope-type calculations. They can all be performed by hand, or using a simple spreadsheet model. This study has shown that the easily calculated the Naïve-Space model works quite well on balance, even though there were simulation cases where the method performed poorly. While the simulated annealing-based "near-optimal" procedure is superior to the naïve methods, using it may not be practical for restaurants with limited resources. Certainly the limiting factor is not the computer time required—it required less than one minute, on average, for each restaurant context to simulate the four naïve table mixes and run the near-optimal heuristic.

The information on customers lost by party size, which is reported here for the first time, may offer insight regarding potential modifications that would improve the naïve methods' performance. Other potentially fruitful areas for future research would consider the effect of rounding rules within the naïve methods and evaluating the use of the best of the four naïve table mixes as a starting point for a SimAnnealSearch-type heuristic, rather than always starting with the Naïve-Space table mix, as I have done here and before (Kimes and Thompson 2005).

A by-product of this simulation experiment was information about how certain design- and operational-related decisions affect revenue. From an operational perspective, these results suggest that a significant revenue benefit can be derived from tighter right-sizing. That is, restaurateurs should maintain a policy of seating parties in tables that fit the party size and are not substantially larger. Fortunately, at least one study has shown that customers view such seating policies favorably (McGuire and Kimes 2006). I also note that the simulations do not necessarily account for customer reactions to a particular policy. For instance, the results showed that using the largest-party table-assignment rule yields higher revenue than applying the longest-waiting-party table assignment rule, but customers may view this approach negatively, as being unfair. Finally, I observed an important benefit of having tables with both odd and even numbers of seats in the mix, compared to tables with just even-numbered seats. This demonstrates the wisdom of having a variety of table types.

Declaration of Conflicting Interest

The author(s) declared no conflicts of interest with respect to the authorship and/or publication of this article.

Funding

The author(s) received no financial support for the research and/or authorship of this article.

References

Bell, Rick, and Patricia L. Pliner. 2003. Time to eat: The relationship between the number of people eating and meal duration in three lunch settings. *Appetite* 41:215-18.

Field, Arthur, Mark McKnew, and Peter Kiessler. 1997. A simulation comparison of buffet restaurants: Applying Monte Carlo modeling. *Cornell Hotel and Restaurant Administration Quarterly* 38 (6): 68-79.

Kimes, Sheryl E. 2004. Restaurant revenue management. A research report of the Center for Hospitality Research at Cornell University, 4(2), Ithaca, NY.

Kimes, Sheryl E., Richard B. Chase, Sunmee Choi, Philip Y. Lee, and Elizabeth N. Ngonzi. 1998. Restaurant revenue management: Applying yield management to the restaurant industry. *Cornell Hotel and Restaurant Administration Quarterly* 39 (3): 32-39.

Kimes, Sheryl E., and Stephani K. A. Robson. 2004. The impact of restaurant table characteristics on meal duration and spending. *Cornell Hotel and Restaurant Administration Quarterly* 45 (4): 333-46.

Kimes, Sheryl E., and Gary M. Thompson. 2004. Restaurant revenue management at Chevys: Determining the best table mix. *Decision Sciences* 35 (3): 371-91.

Kimes, Sheryl E., and Gary M. Thompson. 2005. An evaluation of heuristic methods for determining the best table mix in full-service restaurants. *Journal of Operations Management* 23 (6): 599-617.

McGuire, Kelly, and Sheryl E. Kimes. 2006. The perceived fairness of waitlist management techniques for restaurants. *Cornell Hotel and Restaurant Administration Quarterly* 45 (2): 121-34.

Orkin, Eric B. 1998. Wishful thinking and rocket science: The essential matter of calculating unconstrained demand for revenue management. *Cornell Hotel and Restaurant Administration Quarterly* 39 (4): 15-19.

Thompson, Gary M. 2002. Optimizing a restaurant's seating capacity: Use dedicated or combinable tables? *Cornell Hotel and Restaurant Administration Quarterly* 43 (4): 48-57.

Thompson, Gary M. 2003. Optimizing restaurant table configurations: Specifying combinable tables. *Cornell Hotel and Restaurant Administration Quarterly* 44 (1): 53-60.

Thompson, Gary M. 2007. Restaurant capacity effectiveness: Leaving money on the tables. A research report of the Center for Hospitality Research at Cornell University 7(7), Ithaca, NY.

Thompson, Gary M. 2009. (Mythical) revenue benefits of reducing dining duration in restaurants. *Cornell Hospitality Quarterly* 50 (1): 96-112.

Thompson, Gary M., and Robert J. Kwortnik, Jr. 2008. Pooling restaurant reservations to increase service efficiency. *Journal of Service Research* 10 (4): 335-46.

Thompson, Gary M., and Heeju J. Sohn. 2009. Time- and capacity-based measurement of restaurant revenue. *Cornell Hospitality Quarterly* 50 (4): 520-39.

Bio

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