Computer Simulation

in Hospitality Teaching, Practice, and Research

Computer simulations that test new ideas or procedures can save time, money, and customer inconvenience and dissatisfaction.

BY GARY M. THOMPSON AND ROHIT VERMA

This paper focuses on the use of computer simulation as a tool and how it can be applied in hospitality education, practice, and research. Simulation involves the modeling of an existing real system—for example, skiers' activities at a ski resort,¹ diners' experiences in a full-service restaurant,² and guests' transactions at a hotel's front desk.

Simulation models are used chiefly because it's expensive, complicated, and risky to test proposed new operating procedures in an actual hospitality operation. In contrast, simulation models allow the user to try out different strategies or

alternatives, without actually implementing them in practice. This can reduce costs, since changing real systems can be expensive. Models can also reduce consumer dissatisfaction, since customers, in general, do not like to be part of "failed experiments."³ Once the value of a new idea has been verified through a carefully constructed simulation, that idea may be implemented with a great deal of confidence.

Another reason to build simulation models is that it can be difficult to predict outcomes in complex, real-world hospitality environments, due to the effects of chance. For ex-

¹ Madeline E. Pullman and Gary M. Thompson, "Evaluating Capacity- and Demand-management Decisions at a Ski Resort," *Cornell Hotel and Restaurant Administration Quarterly*, Vol. 45, No. 6 (December 2002), pp. 25–36.

² Gary M. Thompson, "Optimizing a Restaurant's Seating Capacity: Use Dedicated or Combinable Tables?," *Cornell Hotel and Restaurant Administration Quarterly*, Vol. 43, No. 4 (August 2002), pp. 48–57.

³ For example, recall how upset customers became when Amazon.com attempted to institute variable pricing.

Numerical Simulation, Example 1

A Caribbean hotel is considering switching to a self-insurance program to protect itself from the expense of hurricane damage. Assume that there are two categories of hurricane damage: minor, with a probability of 0.02; and major, with a probability of 0.01. The monetary cost of damages for minor and major hurricanes is estimated to be \$500,000 and \$10,000,000, respectively. Thus, the expected annual cost of damages is \$100,000 ($= \$500,000 \times 0.02 + \$9,000,000 \times 0.01 + \0×0.97). What is the relationship between the probability of solvency and the life of the property, the amount set aside for self-insurance, and the maximum sustainable loss?

One can readily model this situation in a spreadsheet. Exhibits 1 and 2 show the results of such modeling efforts. For the case where the maximum tolerable loss is \$1,000,000, Exhibit 1 illustrates the relationship between the probability of solvency, the life of the property, and the amount set aside for self insurance. The chart illustrates that there is actually little effect of even doubling the amount set aside for self-insurance.

For the case where the annual set-aside for self-insurance is \$100,000, Exhibit 2 illustrates the relationship between the probability of solvency, the life of the property, and the maximum tolerable loss. The exhibit illustrates that the maximum tolerable loss has actually little effect on the probability of solvency.¹—G.M.T. and R.V.

¹ The explanation for the counter-intuitive results for the first numerical-simulation example is as follows.

Case 1

- There is little difference across different amounts set aside annually. One would need to build up a reserve of at least \$9,000,000 to survive a major hurricane (the \$10,000,000 damages, less the \$1,000,000 maximum sustainable loss). At a savings rate of \$100,000 per year, it would take 90 years to accumulate sufficient funds. At a savings rate of \$200,000 per year, it would take 45 years to accumulate sufficient funds. Thus, a single major hurricane in the first 45 years of life will cause insolvency, regardless of the amount set aside in the range of \$100,000 to \$200,000 per year. It is only after year 45 that the \$200,000 set aside annually shows any difference, which it does by increasing the likelihood of solvency.

Case 2

- There is little difference across maximum tolerable losses. Here, the amount set aside is \$100,000 per year. With a maximum tolerable loss of \$500,000, it will take 95 years to accumulate sufficient funds to survive a major hurricane (50 years to accumulate sufficient funds with a \$5,000,000 maximum tolerable loss). Thus, any major hurricane in the first 50 years of the property's life will cause insolvency. The advantage of a higher sustainable loss will only become apparent with a property life of more than 50 years.

ample, it may be well known that, in a particular restaurant, 30 percent of its parties are singles and 70 percent of the parties are deuces—but the managers cannot predict with certainty the size of the next party. That's to say, some unpredictability is a key component—and virtue—of simulation models.

Simulation models allow managers to evaluate multiple service designs and perform "what if" types of analyses. Such models can also be used to evaluate the effects of any changes in operating or marketing strategies of hospitality firms (for example, the influence on customer-arrival rates or service requests when specific marketing promotions are introduced). Managers can use computer simulation to determine the best way of controlling the flow of customers and materials, and to find the most effective way to schedule and deploy resources. Simulation replaces the wasteful and often-unreliable practice of testing managers' ideas through trial-and-error methods. As a general guideline, computer simulation is an appropriate tool for process analysis if:

- The system has one or more inter-dependent, random variables;
- The system dynamics are extremely complex; and
- The objective is to observe system behavior over a period of time.

In this article we first discuss the main types of simulation—numerical simulation and discrete-event simulation. We then discuss the tools for building simulation models and offer a few words of caution regarding simulation. We close with our thoughts about the usefulness of computer-simulation modeling.

Numerical Simulation

Numerical simulation involves simulating outcomes controlled by chance, but where the state of the system at specific times is not of interest. The focus in numerical simulations is the aggregate behavior, or outcomes, associated with managers' decisions. These types of simulations are often performed for financial modeling—for example, simulating the effects of particular decisions for a range of chance-controlled outcomes.

Off the shelf. Numerical simulations can typically be performed in computer-spreadsheet software. For example, the function “=RAND()” in Excel® will generate a uniformly distributed pseudo-random number between 0 and 1.⁴ You can then convert the random numbers to serve many other purposes. For example, let’s say we wanted to simulate the size of parties arriving at the restaurant, using the example presented earlier (where 30 percent of parties are singles and 70 percent are deuces). You would first generate a random number, which we call x_p , for party p . We would then use that number, and a translator function, to tell us the size of the party, $size_p$. In this case, our translator function would be:

$$Size_p = \begin{cases} 1, & \text{if } X_p < 0.3 \\ 2, & \text{otherwise} \end{cases}$$

There are standard translation functions to convert the uniformly distributed random numbers to random numbers from other probability distributions, such as normal or negative exponential.

The use of features like Excel’s “Data Tables” enables you to evaluate the outcomes of a chance-based model for a large number of scenarios. You can build a basic spreadsheet model to evaluate one possible outcome (i.e., a “realization”). Using the example in the accompanying sidebars, this would be a sequence of years representing a single possible realization of the lifetime of the hotel. Excel’s Data Table feature (for example) would then be used to simulate a number of possible realizations of the lifetime of the property. By using the Data Table to simulate, say, 1,000 possible lifetimes, one gets a reasonably accurate estimation of

⁴ Random numbers used in simulation are called “pseudo” because they are not purely random. The random numbers are generated using carefully developed functions. Those functions are called “random-number generators.” Random-number generators provide uniformly distributed random numbers between 0 and 1, because such a distribution can readily be converted to many other distributions, and because it is easier to generate random numbers between 0 and 1 than it is to generate random numbers from other distributions.

EXHIBIT 1

Example 1: The probability of solvency as it relates to the life of the property and the amount set aside annually for self insurance, for the case where the maximum tolerable loss is \$1,000,000.

EXHIBIT 2

Example 1: The probability of solvency as it relates to the life of the property and the maximum tolerable loss, for the case where the annual set-aside for self insurance is \$100,000.

Numerical Simulation, Example 2

DataSim¹ is a computer-simulation program designed to facilitate the teaching of managerial issues. The program allows a user to generate potentially any number of constructs, set the characteristics of those constructs, and specify the relationships among the constructs. The constructs may be of many types (e.g., normally distributed, binary, categorical, skewed, truncated), may have varying levels of error in their measurement (i.e., the level of reliability for each measure can be set), and may be approximated through multiple measures of each construct. The user can specify the relationships between the constructs by setting specific correlation coefficients. Once the characteristics of the data set are determined, any number of individual cases can be generated that will exhibit the specified characteristics.

While DataSim can be used for a variety of research and teaching purposes, to date it has been primarily used for teaching. One such example of its use is a sample human-resources staffing exercise. In this exercise, students are given a set of résumés from two hypothetical universities. Students may then apply any number of selection devices (e.g., unstructured interview, structured interview, cognitive-ability test, personality test, job simulation) to any number of the applicants. DataSim was used to generate "scores" for each applicant on every potential selection device. Those scores were based on a wealth of existing research on the validity of the various selection devices and the relationships among the various characteristics being measured. By having specific scores for any number of selection devices for a large number of simulated applicants, students can gain experience on the issues related to the design of selection systems, how information gained from multiple selection devices can be used to make hiring decisions, and on the ultimate accuracy of selection systems. DataSim not only allows the generation of those scores, but it created the data in such a way that it allowed the creator of the exercise to make the relationships among the variables correspond to existing research findings.

DataSim also has a variety of other classroom uses. For example, the ability to generate data with various characteristics, distributions, and relationships has allowed a number of courses on research methods to have richer exercises to teach analytical methods. The program has also allowed teachers to generate hypothetical companies (such as with data on employee-performance levels, job satisfaction, intent to turnover, turnover data, and pay data) that students can then employ in complex class exercises.

In short, the value of the DataSim simulator is not simply what kind of data it can generate, but the fact that such data generation is easy to do, can be done for large numbers of different groups, and can be used to create rich teaching cases that give students a wealth of opportunities to practice the skills they are being taught in the classroom.—G.M.T. and R.V.

¹ Michael C. Sturman, *DataSim*, School of Hotel Administration, Cornell University, 2002.

the probabilities of the outcomes the hotel would experience in its single lifetime.⁵

Numerical simulation can also be used to generate data that has certain characteristics. Michael Sturman, at Cornell University's School of Hotel Administration, has developed such a tool. His tool, DataSim, is used at a number of universities and is described in the sidebar on this page.

Discrete-event Simulation

Discrete-event simulation is applicable where the state of a system over time is the major concern. The term "discrete event" describes the nature of such systems, where the system changes at discrete times when particular events occur. A simple example of a situation where discrete-event simulation could be applied would be the check-in process at a hotel. The events that occur are guest arrival, service commencement, and service completion. Depending on the state of the system when the guest arrives, the customer either commences service immediately or joins a queue. Building a simulation model of the check-in process allows you to evaluate a number of system characteristics, including the effects of new technologies, different staffing levels, and a flexible check-out time.

Discrete-event simulation has three major uses: for validating the results of other models, for modeling alternative-process designs, and for creating management-decision games. We discuss each below.

Validating other models. Simulation is a valuable tool to validate other decision models. Simulation enables you to evaluate theories under a variety of chance-controlled environments. If the theories perform well in the simulated environments, it increases your confidence that the theo-

⁵ Excel's Data Tables are typically used for scenario evaluation. One constructs a table, where the values in the left-most column are repetitively plugged into a cell in the spreadsheet that contains an input variable of interest. Data Table then tabulates the specified "results" cells under the different values of the input parameter. To use Data Table for numerical simulation, the left-most column can be the number of the realization being considered. These numbers are then input into any blank cell in the spreadsheet. If the auto-recalculation feature is turned on, each time Data Table plugs the realization number into the specified blank cell, one obtains a different possible (i.e., simulated) realization.

ries will also perform well in practice. Again, we want to implement only theories or ideas that have a strong likelihood of actually performing well. The sidebar at the top of this page presents an example of how simulation can be used to validate labor-scheduling models.

Process design. Discrete-event simulation is useful for evaluating alternative process designs. Examples of processes for which simulation would be useful are restaurant back-of-house operations and queuing policies at theme parks, ski resorts,⁶ and QSRs. It is uncertainty about customer-arrival times, customer-service times, and—in the case of theme parks and ski resorts—uncertainty about which services customers select, that makes simulation an appropriate and highly effective investigative tool.

Generally, simulation will only identify the best of a number of pre-specified designs. However, in some cases, including the restaurant-table simulation presented in the sidebar on the next two pages, simulation can be embedded into logic that searches for the best possible decision.

Building Simulation Models

Simulation models are generally constructed in spreadsheets, using simulation-modeling tools, or using general-purpose computer languages.

Spreadsheet-based models. As noted earlier, numerical-simulation models can generally be developed in spreadsheets. Modelers with knowledge of Excel's programming language, VBA (Visual Basic for Applications), can build sophisticated discrete-event simulation models (such as the TABLEMIX model mentioned in the sidebar on the next two pages). Advantages of building simulation models in spreadsheets are that knowledge of spreadsheets is widespread and the cost of spreadsheet software is relatively low. Disadvantages, which apply mainly for complex, discrete-event models, include: the knowledge of programming needed, the time required to develop complex models, and the relatively slow speed at which the simulation model executes.

Simulation-modeling tools. Simulation has been used for the past two decades in commercial applications and in classrooms. During the

text continues on page 93

Using Simulation to Validate Labor-scheduling Models

The first author has used simulation extensively to validate labor-scheduling models. Simulation offered a means of mimicking the uncertainty that occurs in real environments. If a labor-scheduling model was tested without addressing information uncertainty, then researchers and managers could question whether the model would, in fact, work well in practice, where information is never perfect.

Here is an example of how simulation was used in this context. The first author developed labor-scheduling models that either (1) deliver a specified level of customer service at the lowest possible cost, or (2) maximize the level of customer service that can be provided at a specific labor cost.¹ The types of uncertainty included in the simulation model to validate the new labor-scheduling models included the arrival rate of customers and the relationship between customers' waiting times and customer satisfaction. Because the models performed better than any existing models, even with the information uncertainty, then one can safely conclude that the models are likely to work better in practice, as well.—G.M.T. and R.V.

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

Management Decision-making Games

Simulation can form the backbone of management decision-making games. Again, the key feature of appropriate decision-making environments is that chance plays an important role. The first author is currently developing two hospitality games for teaching purposes: the Labor Game and the Table Game.

In the Labor Game, one plays the role of a front-desk manager at a large hotel. One has to decide the number of people to schedule and to manage the system in simulated real-time by deciding when to: send employees on break, recall employees from break, send employees home early, or call additional employees in to work. The objective of the game is to provide the best possible service while staying within budget. Chance components of the game include when customers arrive and their length of service and whether employees perform as expected.

In the Table Game, one manages restaurant tables in simulated real time. Chance components of the game include the arrival time, size, and dining duration of parties. Decisions are which party to seat at available tables and whether to hold open tables for expected parties or to later combine with other tables to seat large parties. The objective of the game is to maximize the revenue generated by the tables.—G.M.T.

⁶ Pullman and Thompson, *op. cit.*

Restaurant-table Simulation

The first author developed a model, called TABLEMIX, for simulating restaurant-table occupancy. The model allows one to investigate how different table mixes perform in a restaurant, given information about the restaurant's customer mix.

TABLEMIX inputs include the number of days to simulate, the limit on the number of waiting parties, and the rule used in assigning waiting parties to available tables (e.g., to the party waiting the longest or to the largest waiting party). Inputs for each party size are the probability of that size party, the mean and standard deviation of dining time, the distribution of dining times (normal or log normal), the maximum wait tolerated, and the estimated monetary value of the party. TABLEMIX also requires information on the number of parties expected to arrive, by 15-minute periods, for a peak window of a user-defined duration. If you are simulating a specific restaurant layout, TABLEMIX requires that you specify the position of each table and identify which tables can be combined with which other tables (to seat larger parties).

TABLEMIX can be run in a visual mode, in which case it shows the status of the restaurant—that is, which tables are occupied, how many seats are occupied at each table, which tables have been combined into a larger table, and the number and size of parties waiting for tables (color-coded by the length of time they have been waiting).

Exhibit 3 presents a screen capture of the visual simulator in action (in this case, at the simulated time of 9:28 PM). The main portion of the screen shows the tables. Colored circles represent unoccupied seats, gray circles are occupied seats, while cross-hatched circles represent tables that are “on hold” awaiting the departure of the party so that the table can be combined with an adjacent table to seat a large party. The third table from the left in the bottom row is actually two 4-tops that have been combined to seat a party of five. Those tables will be separated once the party has completed dining. Also shown in the exhibit is a visual display of the waiting parties. Parties that have been waiting the longest appear at the top of the box. As of 9:28 PM, there are eight parties waiting for tables: one party each of 9, 8 and 7 people; two parties

EXHIBIT 3

A screen capture from the visual simulation mode of the TABLEMIX restaurant simulator

of 5 people; and three of 2 people. Finally, the visual mode also reports the number of parties and customers that have been served and lost during the current dining period. As of 9:23 PM, for example, 141 parties of two have been served, and 594 total customers have been served.

TABLEMIX outputs fall into three categories: resource use, customer service, and aggregate measures. TABLEMIX presents the utilization of each resource (i.e., each size of table and seats), by 15-minute time periods, both in tabular and graphic formats. Customer-service measures, which are provided for each party size, include the number of customers served and lost, the expected value of customers served and lost, and the average and maximum waits. Aggregate performance measures are seat use, number of customers served per available seat, revenue per available seat hour, total number of customers served and lost, total expected value of customers served and lost, and the average customer waiting times.

TABLEMIX can be used to evaluate the performance of a specified mix of tables, or it can be used to find the best mix of tables given a specified limit on the number of seats.

Exhibit 4 shows the utilization of seats and tables for a specific mid-scale restaurant, using its existing table mix. Exhibit 5 shows table use under the ideal mix of tables. Seat use is similar in the two exhibits (4 and 5), peaking during the 6:00–8:00 PM time period in the range of 61–63 percent for the existing table mix and in the range of 61–65 percent for the ideal mix. However, the key story is told by the table utilizations.

As Exhibit 4 shows, the current mix of tables is close to maximum capacity, averaging about 95-percent utilization for the four-tops for the 6:00–8:00 PM period, and hitting 90-percent utilization of the 6-tops during the period. By contrast, the ideal mix of tables yields table use of under 75 percent during the peak 6:00–8:00 PM period. The net effect of optimizing the mix of tables is that the effective capacity of the restaurant is increased. The enhanced capacity can enable the restaurant to process about 40-percent more customers during the peak dining period.¹—G.M.T.

¹ Assuming, of course, that the back-of-house workers and the wait staff can handle the increased business.

EXHIBIT 4

Seat and table utilization for the *existing* mix of tables in a specific mid-scale restaurant

EXHIBIT 5

Seat and table utilization for the *ideal* mix of tables in a specific mid-scale restaurant

ServiceModel®

ServiceModel® allows users to design virtually any service process and visually evaluate its performance over time.¹ Users can decide the layout of the service process to be simulated, the customer-arrival rate (including market segments); the number and schedule of service providers; and capacity, resources, and other service attributes. Based on the requirements and model assumptions, the user can redesign various alternatives and run the model for several hours, days, and months of "real" time within minutes of simulation time. In addition, ServiceModel comes with an optimization- and experimental-design program, SimRunner, that can be used to test whether changes in certain input parameters influence selected output measures. SimRunner can conduct full or fractional factorial experiments with multiple attributes with user-defined objective functions. The optimized regression models can estimate both the main effects and interaction effects among input variables on the objectives. The estimated equations can be easily incorporated into spreadsheet-based decision-support systems for use in day-to-day

management and decision making. In fact, the most recent version of ServiceModel is compatible with Microsoft's ActiveX control protocols; that allows users to control various functionalities of the simulation via Microsoft Excel macros.

The illustration above shows a screen shot of the simulation of customer orders arriving at a restaurant kitchen.² During peak periods the kitchen becomes overloaded and orders queue up, adversely affecting customers' wait times. The model tells us where the bottlenecks are by showing us the work loads of the cooks and the various pieces of kitchen equipment. It also reports queue times, lengths, orders processed per hour, and customer waiting times. Different managerial approaches can be tested using the model to see which ones are effective in solving the bottleneck conditions. Using the model, a cost-benefit analysis can be performed to see which solution provides the most benefit for the least cost.—G.M.T. and R.V.

¹ See: www.servicemodel.com (as viewed on 8-27-02).

² ServiceModel2001 Demo, www.servicemodel.com (as viewed on 8-27-02).

continued from page 89

DOS and mainframe-computing era, programming languages (for example, FORTRAN, C, C++) were used to develop simulation models. Although sophisticated and detailed, those programs had either limited or no ability to graphically display the models. Additionally, programming even a simple model required several hours (often months) of development time.

More-recent simulation programs, on the other hand, are relatively easy to use, display information graphically, and do not require knowledge of a programming language. They do, however, require an understanding of simulation-modeling concepts, logic, and statistics. A number of graphics- and animation-based simulation programs have been developed that attempt to reduce model-development time. XCELL[®] is an excellent example of one of the first widely used graphics-based simulation-modeling programs.⁷ Currently there are a number of animation-based, simulation modeling tools available in the market.⁸ The sidebar on the previous page describes the capabilities of one leading simulation program, ServiceModel[®], which the second author has used extensively in teaching, research, and consulting applications over the past several years.⁹

Advantages of simulation-modeling tools are that one does not require knowledge of programming to create complex discrete-event models, and the fact that models can be created quite quickly by experienced modelers. The speed with which simulation models can be constructed with these tools make the tools suited to rapid prototyping of process designs. Disadvantages of simulation-modeling tools include a slow speed of execution and the cost of the software, which can be in the tens of thousands of dollars per copy.

⁷ R. Conway, W.L. Maxwell, J.D. McClain, and S.L. Worona, *XCELL Factory Modeling System*, Version 4.0, third edition, Scientific Press, San Francisco, California, 1990.

⁸ For reviews and comparisons of various commercially available simulation tools, see: <http://lionhrtpub.com/orms/surveys/Simulation/Simulation.html>

⁹ R. Verma, G. Gibbs, and R. Gilgan, "Redesigning Check-processing Operations Using Animated Computer Simulation," *Business Process Management Journal*, Vol. 6, No 1 (2000), pp. 54-64.

Words of Caution

Simulation, like any quantitative decision aid, is dependent on accurate data. If a simulation model is created that has incorrect or inaccurate parameters, then the results of that model will be incorrect. Certainly the old adage "garbage in, garbage out" applies to simulation. To overcome this liability, it is important to validate a simulation model. One can use a variety of means to do that. First, the model can be used to simulate an existing system. The behavior of the model should match that of the existing system. Second, one can ask experts in the area the type of behavior they would expect the system to exhibit under certain situations, which can then be compared to the simulation model's outputs under those circumstances. Such experts are then giving the simulation model a sort of "reality check."

Another potential drawback of simulation relates to its ease of use. Because simulation models are relatively easy to create, simulation is often overused. In other words, simulation is used when other tools, such as mathematical programming, might be more appropriate. The way this limitation can be overcome is by ensuring that the modelers have a broad knowledge of different decision aids, and that they know which tools are most useful for which types of situations. The point here is that if all one knows is simulation then everything tends to look like a simulation problem.

Conclusion

Simulation, when used correctly, is a useful tool in hospitality contexts. It can be used in practice, in designing better customer-service processes, in research,¹⁰ in validating new theories (like the validation of the labor-scheduling models), and it certainly belongs as part of a hospitality curriculum (like the numerical-simulation examples and the management decision-making games). Simulation is typically used for identifying the best of a number of predefined alternatives, in cases where the effects of chance make it difficult to perform other analyses. Our recommendation is to "take a chance—with simulation!" ■

¹⁰ See, for examples: Thompson, *op. cit.*, and Verma, *op. cit.*

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

Rohit Verma, Ph.D. (bottom photo), is an associate professor of operations management at the David Eccles School of Business, University of Utah, Salt Lake City (rohit.verma@business.utah.edu).