

SPREADSHEET SIMULATION

Andrew F. Seila

Terry College of Business
The University of Georgia
Athens, GA 30602, U.S.A.

ABSTRACT

“Spreadsheet simulation” refers to the use of a spreadsheet as a platform for representing simulation models and performing simulation experiments. This tutorial explains the reasons for using this platform for simulation, discusses why this is frequently an efficient way to build simulation models and execute them, describes how to setup a spreadsheet simulation, and finally examines some limitations on the use of spreadsheets for simulation.

1 INTRODUCTION

A simulation is a sampling experiment that is done on the computer (Fishman 1996). At the core of any simulation is a model that involves quantities whose values are unpredictable and therefore must be sampled from an appropriate population. The model is represented using a computer program, and the program actually samples the random variables, performs the computations of the model and reports the outcome, usually in the form of one or more numerical values. All simulation models fit this description. *Spreadsheet simulation* simply involves the use of a spreadsheet to represent the model, do the sampling, perform the model computations and report the results.

The idea of an electronic spreadsheet for storing information and performing calculations dates back to Mattesich (1961) when computers were mainframes, but spreadsheets as we know them are a product of the microcomputer age. The earliest spreadsheet, VisiCalc, had limited functionality but introduced the world to the concept of the interactive electronic spreadsheet, i.e., a program that stores data and computations in a rectangular array of cells and allows the user to manipulate the values and formulas interactively. Later spreadsheets such as Lotus 1-2-3™, Microsoft Excel™ and Quattro Pro™ greatly expanded the features and developed the spreadsheet into an effective modeling, prototyping, analysis and presentation tool.

Today, spreadsheets are available for all of the major desktop operating systems: Windows, Unix/Linux and Mac OS. Since Mac OS X is a derivative of BSD Unix, all of the Unix spreadsheets (as well as other software) will eventually be available for this platform. The most prevalent spreadsheet today is Microsoft Excel™, which is part of Microsoft Office™. Most other spreadsheets operate similarly and have similar features as Excel™. In this paper, we will refer to Excel™ menus and other features. However, the concepts and techniques to be discussed apply to virtually all other spreadsheets.

Spreadsheets are used by many people. Since the idea of the spreadsheet started in the context of accounting, it is natural that they are widely used in business. Most business users, however, employ a very small subset of the available features. Engineers have been slow to adopt spreadsheets as a computation and analysis platform, perhaps because they have been trained to use other software tools such as MATLAB™ that were written specifically for mathematical modeling. These tools are very powerful and can certainly be used to develop a broad range of simulation models, but in many cases a spreadsheet is simpler and more intuitive to use.

2 WHY USE A SPREADSHEET FOR SIMULATION MODELING?

We will use the term *simulation platform* to refer to the software environment used to develop, test and run a simulation experiment. First, let’s examine what features a simulation platform needs. Following is a list of the capabilities that must be available:

1. A way to represent mathematical and logical relationships between variables in the form of computations and assignment of values, and algorithms that describe how to do a series of computations.

2. A way to generate uniformly distributed pseudorandom numbers and use them to sample observations from various distributions.
3. A means to repeat the computations of the model, thus implementing replications.

This list is minimal, but these are the necessary capabilities for the platform to be used for simulation. Most spreadsheets have these features. The following useful features are also available in most spreadsheets to make the process quick and reliable:

1. A large number of functions to do mathematical, statistical, database, date/time, financial and other calculations.
2. Database representations and database access.
3. Charting and graphing.
4. Display and documentation features such as fonts, colors and geometric shapes to improve presentation.
5. Automation through scripting languages such as VBA (in the case of Excel™).

The table structure of spreadsheets allows the developer to organize the computations and outputs in a natural and intuitive manner. Spreadsheets are ubiquitous—almost everybody has one—and file formats are standardized, so files written by one spreadsheet can be imported by others. Spreadsheets have become a standardized platform to represent computations. As a result, developers and users can easily pass simulation models from one to another. For these reasons, the spreadsheet is an attractive platform for simulation.

There are a number of publications that discuss spreadsheet simulation. See Winston (1996) and Seila et al. (2003), Chapters 2 through 4 for very readable tutorials.

3 WHEN IS SIMULATION USING A SPREADSHEET APPROPRIATE?

Certain modeling situations lend themselves well to implementation in a spreadsheet. Indeed, any set of calculations in a spreadsheet can be considered a model. Usually, these models have parameters or variables whose true values are unknown and thus assumed values are used.

3.1 Stochastic Models

In some cases, unknown parameters such as the interest rate at a future time or the demand for a product are actually random variables whose value cannot be predicted, i.e., the models are stochastic. Many stochastic models in finance (including real estate and insurance), logistics and

engineering can be conveniently setup in a spreadsheet for simulation. Spreadsheets are frequently used by actuaries, for example, to evaluate insurance rating methods. Consider, for example, an inventory model in which the demand for the product is stochastic. In order to evaluate a particular replenishment policy, this value must be sampled when the simulation experiment is run. An experiment would consist of sampling demand for the product and applying the inventory policy over a long period of time to compute observations of the periodic costs resulting from excess inventory and shortages associated with the policy. These observations would then be used to estimate the mean cost for the policy. The experiment would be repeated for several policies to find the inventory policy that produces the minimum mean cost. This is a typical stochastic model that can be analyzed using simulation and will be used as an example later in this paper. Interesting spreadsheet implementations of queueing simulations have also been developed (Grossman 1999).

3.2 Sensitivity Analysis for Spreadsheet Models

Another situation where spreadsheet simulation is useful involves doing a “what-if” or sensitivity analysis for models having parameters that are unknown but not necessarily random. It is often the case in spreadsheet models that modelers want to determine how sensitive the performance measure is to variations in these parameters. For example, in a model that concerns the leasing or purchase of a piece of real estate, the mortgage interest rate at the time the contract is signed is an unknown parameter. The present value of each decision (lease vs. buy) will depend upon this parameter value and the decision maker would like to know how sensitive the present value of a policy is to it. If only one or two parameters are involved, modelers can use the “Table” command of the “Data” menu to evaluate the performance measure when each parameter value in a collection of possible values is substituted into the model. Excel™ and other spreadsheets will support this calculation with one unknown parameter and many performance measures, or with two unknown parameters and one performance measure. For example, one could vary the interest rate from 5.5% to 11.0% in steps of .5% and use each value to compute the present value of the lease decision and the present value of the buy decision. Or, one could vary the interest rate and also vary the value of the property over a discrete set of values, and for each combination of these two values, compute the difference between the present values of the two decisions.

Real spreadsheet models normally have many unknown parameters, as well as multiple performance measures. This type of what-if analysis can become unwieldy when the model has more than a few parameters. For example, suppose that the number of unknown parameters in the

model is 10, and the number of possible values for each of these parameters is 3, denoting the minimum, most likely and maximum values. Then, the number of recalculations that must be performed in order to assess all combinations of these possible values is $3^{10} = 59,049$. Clearly, this is possible only if the process of recalculation is automated, and then it is rather time-consuming. If the number of parameters increases to just 15, the number of recalculations grows to about 14 million, an infeasible computational task on most desktop systems. The solution to this conundrum is simulation. By sampling these unknown values from appropriate distributions, one can do a “what-if” analysis on a model with a large number of unknown parameters. In fact, 1000 replications generally produces enough observations to assess the variation in the output measures, regardless of the number of combinations of values of unknown parameters. Thus, simulation is a useful technique when the number of unknown parameters is moderate or large.

The mechanics of setting up and running a spreadsheet simulation are very much the same in both of these cases, but there is one important difference in the way the output data are analyzed: When simulating a stochastic model, you are usually interested in using the output data to estimate an unknown performance measure for the model; when doing sensitivity analysis, you are interested in using the output data to assess the amount of variation in one or more output quantities.

4 HOW DOES ONE SETUP A SPREADSHEET SIMULATION?

The cells in a spreadsheet model can be classified by their contents:

- *Inputs to the model.* These cells can contain parameters that are part of the model, such as unit costs or mean demand. The contents can also be sampled values of the random variables that represent uncertain quantities in the model such as demand or price paid, or they can be assumed values of unknown parameters when one is doing a sensitivity analysis.
- *Intermediate computations.* These cells contain formulas that have the calculations that define the model. For example, in an inventory model, they might compute the inventory levels or backlogs at the end of each period. These computations define the transformations that convert the model’s inputs to the outputs.
- *Outputs from the model.* These cells contain the observations on quantities of interest one seeks from the model. For example, in an inventory model, these observations could be the costs incurred during each period.

Most models that can be organized in this way can be simulated in a spreadsheet. The following steps are described in more detail in Chapter 2 of Seila et al. (2003). A simple inventory model will be used to illustrate the process.

4.1 An Inventory Model

Consider a single period inventory model where a quantity of a good will be purchased to satisfy a stochastic demand whose distribution is known. As an example, we could be placing an order for the number of hot dogs for a baseball game. Demand will be determined by many unpredictable factors, but data from past games shows that it has an Erlang distribution with parameters 4.0 and 2. This random variable has mean 8.0. We experience costs of $c_e = \$60$ per case for an excess (if the amount ordered is greater than demand) and $c_s = \$160$ per case for a shortage (if the amount on hand is less than the demand). Let D represent the demand—a random variable—and x represent the number of cases ordered—a decision variable. We can order cases in fractional amounts. Then, the realized cost, after we have attempted to satisfy demand, is

$$Y = \begin{cases} c_e(x - D) & \text{if } x > D, \\ c_s(D - x) & \text{if } x \leq D. \end{cases}$$

We want to simulate this model in order to estimate the expected cost, given a specific order amount, x . If we do this for several values of x , we can select the order amount that provides the minimum cost.

4.2 Setup the Model

The first step is to build the model in the spreadsheet using fixed values for all parameters and other inputs, including those that will be sampled. This lets us check the computations and assure the correctness of the model transformations before the simulation-specific components are added. In our example, we would create the model with a specific value for demand, which is the only stochastic quantity in the model. Figure 1 shows a portion of the spreadsheet containing the model at this stage. Cells D4 and D5 contain the cost parameters, c_e and c_s , respectively. Cell D7 has the order amount, which is 9 in this case. Cells D9 and E9 have the two parameters for the distribution of demand—the mean and “stages” parameter for the Erlang distribution. The value in cell E9 is not used at the moment, but we entered it now so it all parameters will be available at the next step. Cell D12 contains the (assumed) demand, which we set to the value in cell D9 for the moment. Then, the cost is computed in cell D11 using the formula

Figure 1: Basic Inventory Model

$$=IF(D12 < \$D\$7, \$D\$4 * (\$D\$7 - D12), \\ =\$D\$5 * (D12 - \$D\$7)).$$

We can put values in cell D7 that are greater than and less than the values in cell D9 to check the computations and assure that they are correct.

Second, add the random variate generators, i.e., replace the values in the cells that represent sampled quantities with formulas that sample these values from appropriate distributions. Appropriate formulas can be found in any reference on random variate generation. See Cheng (1998) for example. Spreadsheet add-ins for simulation, which we will discuss in a later section, also include formulas for sampling from various distributions. At this point, all random variates can be resampled when the spreadsheet is recalculated. In the inventory example in Figure 1, the formula in cell D12, which contains the demand, is replaced by the formula

$$=\$D\$9 / 2 * dGammaDev(\$E\$9)$$

to sample demand from an Erlang distribution. The specific function used here, `dGammaDev()`, comes from the PopTools add-in, which will be discussed in Section 5. Each time the spreadsheet is recalculated by pressing the F9 key, the contents of cell D12, the demand, is resampled and thus changes. The new value of demand is then used to compute the cost in cell E12. Thus, each press of the F9 key produces a new replication for the simulation, and therefore, a new observation of the cost in cell E12.

Third, identify the “output data” for the model. Actually, the modeler should know these desired performance measures when the model is created. For example, in an inventory model, you might use the mean cost per period as a performance measure, so the output data for the model would be the costs incurred in each period. Here, you want to identify those cells that contain the values of these performance measures. At this point, you can watch the values of these cells change (i.e., being sampled) each

time the spreadsheet is recalculated. Since our example is a single-period inventory model and we want to estimate the the mean cost, the output data is just the contents of cell E12, the observed cost for the replication. Note that in more elaborate models, you could have several performance measures.

4.3 Create the Simulation Run

It is useful to distinguish two types of simulation experiments: (1) static simulations that are run in a series of independent, identical replications and produce independent, identically distributed output observations, and (2) dynamic simulations that are run in one long replication and produce a time series of dependent observations. The setup is different for each of these.

Where independent replications are performed, the model computations are usually contained in some region or group of regions of the spreadsheet. Since a recalculation produces a replication, we need to do a series of recalculations of the spreadsheet and save the outputs after each recalculation to perform the replications. There are several ways to accomplish this. If the model computations can be placed in a single row, we can just copy this row the appropriate number of times and all replications are displayed at once.

Figure 2 shows a portion of the inventory model with replications added in columns C, D and E. These were created by simply entering the first replication number in cell C12, then copying the range C12:E12 down an additional 199 rows to create 200 replications.

	A	B	C	D	E	F
1	Soft Drink Inventory Example - Stochastic Model					
2						
3	Problem Parameters:					
4	Cost per unit for excess			60	Input parameters	
5	Cost per unit for shortage			160		
6						
7	Order Amount			9	Decision variable	
8						
9	Demand distribution parameters		8	2	Inputs	
10						
11			Replication	Sampled Demand	Sampled Cost	Outputs
12			1	19.15	1623.72	
13			2	13.10	656.62	
14			3	1.46	452.68	
15			4	10.86	297.18	
16			5	10.55	248.06	
17			6	8.65	20.74	
18			7	4.45	272.85	
19			8	1.97	422.10	

Figure 2: Inventory Model with Replications

If the computations in the model are more involved and cannot be placed on a single row, we can use the Table command in the Data menu to tell the spreadsheet to go through an iterative recalculation, storing the values of the outputs after each recalculation. To use the Data-Table

command, you must first create a rectangular range of cells to hold the model outputs. Start by creating a column of numbers having values from 1 to the number of replications you will perform. Excel™ and most other spreadsheets have an easy way to create a column or row of consecutive numbers. Above each adjacent column to the right, place a formula that will produce the value of a specific output. The design is for each row of this table to contain the replication number in the first column and the outputs, i.e., observations on each performance measure, for that replication in the adjacent columns.

To actually run the replications, select the Data menu, then select the Table option. This command was originally created to perform “what-if” scenarios as described above by substituting each value in the first column into a specific cell, recalculating the spreadsheet, and recording the values of other cells that depend upon the substituted value adjacent to the substituted value. In our case, the contents of each cell in the first column of the data table is just a replication number, and since the replications are independent and identically distributed, the replication number is not actually used in recalculation. However, the recalculation will cause all random variate sampling formulas to re-execute, producing a new observation for each random variate and thus new observations for all outputs that are statistically independent of those for all other replications. The Data-Table command presents a dialog asking where you want to put the input value. Click in the field labelled “Column Input Cell” and select any unused cell, then click “OK”. The data table containing the replications will fill quickly. When it is finished, each column of this table except the left-most will contain all of the observations on a specific performance measure. you might need to press the Recalculate key, F9, to compute the new values in the table. It is not difficult to do thousands, or even tens of thousands, of replications in this way.

In a dynamic simulation, the output values are observed periodically over time or in sequence. For example, in an inventory model, the costs incurred might be observed at the end of each week. In addition, each output observation will depend in some way on the previous outputs. If each period’s computations can be placed in a single row, then the next period’s computations are constructed from the contents of the cells in the previous row. Once the computations for a representative set of periods are setup, i.e., once a representative row is entered, the row(s) can be copied, thus extending the time span of the model and producing the desired number of periodic observations. As a result, the sequence of dependent output observations in the simulation will be contained in one or more columns of the spreadsheet. Chapter 4 of Seila et al. (2003) has some examples of dynamic models implemented in a spreadsheet.

4.4 Analyze the Data

Each simulation has its own data analysis requirements (Alexopoulos and Seila 1998). For stochastic models, analysis normally involves applying statistical procedures to compute estimates of population parameters as well as confidence intervals for these estimates. When the objective is sensitivity analysis, the data analysis seeks to determine the likely *range of values* of the output data. This can involve computing extreme values of the data such as quantiles and graphically displaying the distribution of the data. Most spreadsheets have formulas for computing the sample mean, sample variance and quantiles of well known distributions such as the normal distribution, so the usual confidence interval formulas can be applied. Spreadsheets also have a rich selection of other statistical computations such as regression analysis and quantile computation, which can be applied too. The tutorial on output data analysis in this Proceedings can provide additional information about applying statistical procedures to simulation data.

Figure 3 shows the inventory example with a simple data analysis in the range G2:H9. In this case, we are just computing a confidence interval for the mean. All formulas used in H2:H9 are standard Excel statistical formulas.

Replication	Sampled Demand	Sampled Cost	Outputs
1	11.06	328.91	
2	17.51	1362.10	
3	13.35	696.45	
4	8.89	6.63	
5	6.24	165.80	
6	2.78	373.21	
7	0.41	515.42	

Figure 3: Inventory Model with Data Analysis

For the single period inventory model which we are using as an example, our objective is to find the order quantity that minimizes the mean cost. Finding this value will require that we repeat the entire simulation experiment using various values for the order quantity, which is in cell D7. This is a good place to use the Data-Table command. The result is shown in Figure 4, in the range G11:I25. First, we created the sequence of order amount values in the range G13:G25. Then, we placed formulas =H5 and =H7 in cells H12 and I12, respectively, to copy the estimates of the mean and sampling error of the mean to these cells. Finally, we selected the range G12:H25 and invoked the Table command of the Data menu to produce the Table dialog in Figure 4. In the field labelled “Column Input Cell”, we entered D7 to indicate that each value in the range G12:G25 should be

placed into D7 before recalculating the spreadsheet. Upon clicking OK in the Table dialog, we get a table resembling that in Figure 4. Since this is a simulation, each spreadsheet will show different values in this table. We can examine the values in this table, but they are much easier to interpret if we display them graphically.

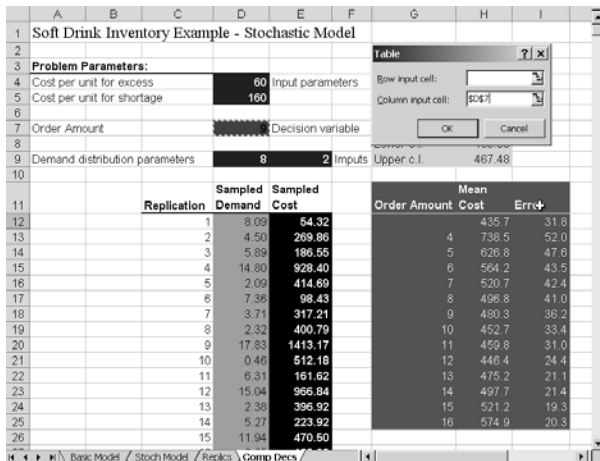


Figure 4: Order Quantity Evaluation

Presentation generally includes some tables and graphs. Spreadsheets have extensive facilities that make it easy to produce these types of presentations in high quality. The types of graphs or other displays will, of course, depend upon the data analysis and the objectives of the modeling effort. Figure 5 shows a graph of the results in Figure 4. From this graph, you can not only identify the order quantity that minimizes cost, but also you can see that the cost is not very sensitive to the order quantity when it is close to the optimal value.

It is important to note that in this example, we have implemented a simulation model, run replications, collected output data, analyzed the output data to estimate the required parameters of the model and used the model to automate the evaluation of a series of decisions. With the exception of the formula for generating random variates, the entire process was done using built-in Excel™ features.

This and some additional spreadsheet simulation models can be found at <http://seila.terry.uga.edu/spreadsheetsSim>. These models illustrate the concepts and techniques just discussed and demonstrate how to implement and run dynamic simulations using a spreadsheet.

5 SIMULATION ADD-INS FOR SPREADSHEETS

The process of developing and running a simulation in a spreadsheet can be simplified somewhat by using one of the available add-in packages for Excel™. @RISK™ <http://www.palisade.com> and Crystal Ball™ <http://www.decisioneering.com> are com-

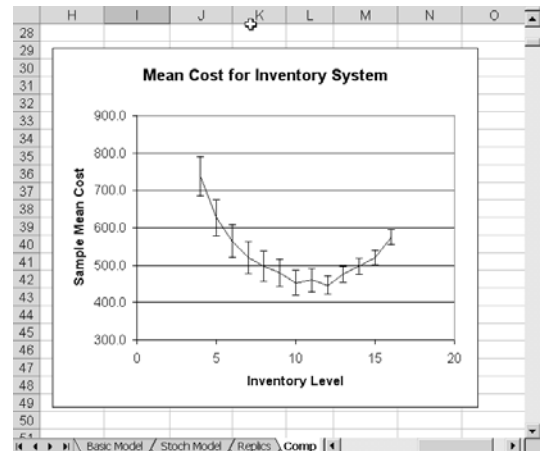


Figure 5: Graph of Expected Cost Versus Order Quantity

mercial packages. PopTools <http://www.cse.csiro.au/poptools/> is a free Excel™ add-in. Another free add-in for Excel™ called SIM-TOOLS.XLA by Professor Roger Myerson is available at <http://home.uchicago.edu/~rmyerson>. These packages provide several features that are not included in the basic spreadsheet:

- Random number generation using documented and tested algorithms.
- Extensive functions for generating random variates from a variety of distributions.
- Features to automate the setup and running of the simulation experiment.
- Features to automate analysis and presentation of the output data from the simulation experiment.
- Optimization procedures for the model.

The random number function, which is called “RAND()” in most spreadsheets, produces a pseudo random sample from a uniform distribution between 0 and 1. Unfortunately, many spreadsheet publishers do not document the algorithm used in RAND(). Frequently, these are just the functions that are distributed as part of the C or C++ compiler. Research has shown that some algorithms for generating random numbers have better statistical properties than others (Fishman 1996, L’Ecuyer 1998). Thus, using the built-in RAND() function carries some risk that the random numbers will not behave as truly independent, random numbers. In @RISK™, Crystal Ball™ and PopTools, the random number generators have been tested and documented, and therefore are recommended over RAND().

It is easy to write functions that generate observations from some distributions such as the triangular, exponential and normal distributions, starting from independent uniform random variates (Cheng 1998). However, observations from some distributions such as the Gamma and Weibull are difficult or impossible to generate using just the built-in

functions of the spreadsheet. These add-ins provide easy, intuitive functions for all common distributions.

If you use the Data-Table method described above to run replications, some effort is required to set it up and the method uses space in the spreadsheet to store the results. These packages implement their own iterative procedure to run replications and store the resulting summary statistics or raw data. Often, you do not need to store all of the raw data. Only the summary statistics are needed. Thus, these add-ins can simplify the problems of setting up and running simulations, and analyzing the output. Examples of the use of these add-ins can be found on their websites and examples of the use of @RISK™ for financial modeling can be found in Chapter 3 of Seila et al. (2003).

6 WHEN IS SPREADSHEET SIMULATION NOT ADVISED?

Spreadsheets are powerful, convenient tools for simulation modeling, but they do have four important limitations.

1. Only simple data structures are available in spreadsheets. The spreadsheet consists of a group of pages, or worksheets, each of which has a table consisting of rows and columns of cells. Each cell can contain data or a formula. One can treat a column or row of cells as a vector, and a two-dimensional range of cells can be treated as a two-dimensional array, or matrix. In some simulation models, more elaborate data structures such as lists and trees are needed. One case in point is that of discrete event simulation, where lists are needed for the event list and waiting lines. These structures can be built in a spreadsheet, but they are contrived and inefficient.
2. Complex algorithms are difficult to implement. For the most part, formulas in spreadsheet cells are static computations that are executed once when the cell is recalculated. Spreadsheets do not have convenient facilities to implement a while-loop or a for-loop. These can also be implemented, but the implementation is often inefficient and inflexible. For example, if a computation needs to be done 10 times, it can be implemented in a column or row of 10 cells. But, what if it needs to be done 100,000 times? Most spreadsheets do not allow a column this long. Moreover, how would you implement a loop that must be executed until a particular value is obtained? For example, in an actuarial ruin model, the value of the firm is computed until it becomes negative. Since you do not know the maximum number of periods to guarantee ruin (it might even be infinite), you do not know how many cells to include. VBA in Excel™ can be used to implement more complex logic, but this is a more advanced tool that is seldom used by casual spreadsheet users. So, models that require complex loops and other conditional computations may not work well in a spreadsheet.
3. Spreadsheets are slower than some alternatives. Consider what a spreadsheet must do to recalculate. Formulas are stored in “source code.” That is, the spreadsheet must interpret the formula before it can be executed. This interpretation action normally takes much longer than the execution. Some spreadsheets are sophisticated enough to store the executable code so the interpretation does not have to be repeated each time, but it is nevertheless a much less efficient setup than one would have with a compiled language. Moreover, spreadsheets use much more of the computer’s resources to support the elaborate user interface and provide all of the features. Thus, spreadsheets are inefficient in their use of memory. A model that is very large and/or requires long simulation runs would need to be programmed in a compiled language in order to execute in a feasible length of time or use a reasonable amount of main memory.
4. Data storage is limited. Since the output data must be stored in the spreadsheet, usually in a column, the length of the output series is limited by the maximum column length. In many spreadsheets, column lengths can be tens of thousands or even hundreds of thousands of cells. However, some models such as those that evaluate the reliability of highly reliable systems, require very large sample sizes - in the millions of observations. There are ways to circumvent this restriction. One could use multiple columns to store output data for the same performance measure, or one could accumulate sample statistics without actually storing the raw data. All of these solutions require a more complex approach to the simulation and result in more inefficiency in the execution of the simulation. When this is the case, it is appropriate to ask if another platform would be a better solution.

These four limitations seem to restrict considerably the range of models that can be implemented in a spreadsheet. However, many models are not subject to these restrictions, and they are often done to get “quick and dirty” results. This can be the case in many business models where a manager or system engineer wants to estimate a parameter to plus-or-minus 10 percent, for example. This is the place where a spreadsheet really earns its bars. Prototypes can be quickly built and run in a spreadsheet. If the prototype shows that the simulation does not work well in the spreadsheet, then it can be moved to a more appropriate platform.

7 CONCLUSIONS

Simulation modeling is used in only a very small percentage of situations where it can provide valuable information for decision making. The reasons for this underutilization are many. Sometimes, the managers or analysts are not familiar with specialized simulation software. Spreadsheets provide an easily-used platform for simulation that is already on the desktop of all analysts, engineers, managers, administrators and others who need to model and simulate problems. Since spreadsheets have powerful functions for doing sophisticated computations and excellent graphing features for displaying the results, they can be used in the entire process from analyzing input data to developing the model to analyzing and presenting the simulation results.

If the model is already implemented in a spreadsheet, managers can experiment with the model and evaluate alternative without having to involve simulation specialists. Many simulations do not need to be extensive. They are designed to provide ball-park estimates and to show general system behavior. This is often true of financial models. These models can usually be implemented most efficiently in a spreadsheet. Simulation problems for which spreadsheets are a useful platform also include prototype models which are relatively small and designed to provide a proof of concept. Simulation is also a useful tool to do a sensitivity analysis for any spreadsheet model that has parameters that are unknown or whose values can change.

Commercial and free spreadsheets are continuing to be developed. Future versions will undoubtedly allow larger worksheets and perform computations more efficiently. As computing power continues to grow, the limitations to spreadsheet simulation will be removed and this platform will be even more attractive. Excel™ comes bundled with an optimization tool (solver). Perhaps there will be a time when it also comes bundled with a simulation tool!

REFERENCES

- Alexopoulos, C., and A. F. Seila. 1998. Output data analysis. In *Handbook of Simulation: Principles, Methodology, Advances, Applications, and Practice*, ed. J. Banks. New York: John Wiley.
- Cheng, R. C. H. 1998. Random variate generation. In *Handbook of Simulation: Principles, Methodology, Advances, Applications, and Practice*, ed. J. Banks. New York: John Wiley.
- Fishman, G. S. 1996. *Monte carlo concepts, algorithms and applications*. New York: Springer.
- Grossman, T. A. 1999. Spreadsheet modeling and simulation improves understanding of queues. *Interfaces* 29 (3): 99–103.
- L'Ecuyer, P. 1998. Random number generation. In *Handbook of Simulation: Principles, Methodology, Advances, Applications, and Practice*, ed. J. Banks. New York: John Wiley.
- Mattesich, R. 1961. Budgeting models and system simulation. *The Accounting Review* 36:384–397.
- Seila, A. F., V. Ceric, and P. Tadikamalla. 2003. *Applied simulation modeling*. Belmont, California: Brooks-Cole.
- Winston, W. L. 1996. *Simulation modeling using @risk*. Belmont, California: Duxbury.

AUTHOR BIOGRAPHY

ANDREW F. SEILA is a Professor in the Department of Management Information Systems in the Terry College of Business at the University of Georgia. He has been involved with simulation teaching and research for more than 30 years, and has attended the Winter Simulation Conference since 1977, serving as Program Chair for the 1994 meeting in Orlando. Dr. Seila is a Fulbright Fellow and recognized authority on the use of spreadsheets for model representation and simulation. His professional interests include statistical methodology for simulation, modeling methodology and simulation applications, especially those in healthcare and finance. He is the author of over 50 refereed papers and co-authored the chapter on output analysis in the *Handbook of Simulation*. He has been a consultant to numerous businesses, consulting firms and public institutions. In recent years, Dr. Seila has expanded his interests to Internet Technology. He is currently Director of the Internet Technology Program at the University of Georgia. His e-mail address is andy@ms.terry.uga.edu.