

SPREADSHEET SIMULATION

Andrew F. Seila

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

ABSTRACT

“Spreadsheet simulation” refers to the use of a spreadsheet as a platform for representing simulation models and performing the simulation experiment. 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 when a spreadsheet is not an appropriate platform for simulation.

1 INTRODUCTION

A simulation is a sampling experiment that is performed on a computer (Fishman 1996). Every simulation starts with a model that includes variables whose values can vary randomly, and therefore, whose values are unknown and 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 results, usually in the form of one or more observations. All simulation models fit this description. “Spreadsheet simulation” simply involves the use of a spreadsheet to represent the model and conduct the sampling process.

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 electronic spreadsheet, i.e., a program that stores data and computations in a rectangular array of cells. 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 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 comes bundled with Microsoft Office. Most other spreadsheets operate similarly and have similar features, so the terminology used in this paper will refer to Excel, while 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, use a very small subset of the available features. Engineering has been slow to adopt spreadsheets as a computation and analysis platform. Engineers have been trained to use other software tools such as MATLAB that were written specifically for mathematical modeling. These tools are 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. One of the goals of this tutorial is to introduce engineers to the spreadsheet as a tool for simulation modeling.

2 WHY USE A SPREADSHEET FOR SIMULATION MODELING?

The term “simulation platform” refers 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 random numbers and sample observations from various populations.
3. A means to repeat a series of identical computations, thus implementing replications.

This list is minimal. If any of these features is missing, the platform cannot be used for simulation. All popular spreadsheets have these features. In addition, the following features are 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 results in a natural and intuitive manner. Spreadsheets are ubiquitous - almost everybody has one - and files written by one spreadsheet can usually be imported by others. 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 Ragsdale (2000) for very readable tutorials.

3 WHEN IS SIMULATION USING A SPREADSHEET APPROPRIATE?

Certain modeling situations lend themselves well to implementation in a spreadsheet. Indeed, every set of calculations in a spreadsheet can be considered a model. Usually, parameters or variables are present in these models that are unknown and used with assumed values.

3.1 Stochastic Models

In some cases, the unknown parameters are actually random variables whose value cannot be predicted, i.e., the models are stochastic models. Many stochastic models in finance (including real estate and insurance), logistics and engineering can be conveniently setup in a 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 short-

ages 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 lowest mean cost. This is a typical stochastic model that can be analyzed using simulation. 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 unknown parameters that are not necessarily random, but simply unknown at the time the model is used. Modelers frequently 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 a parameter that is unknown. The present value of each decision (lease vs. buy) will depend upon this parameter value. This analysis is often called “what-if analysis.” If only one or two parameters are involved, modelers can use the “data table” feature of the spreadsheet that lets one setup a list of values for the unknown parameter and evaluate the performance measure when each parameter value 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, for each value, compute the present value of the lease decision and the present value of the buy decision. 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.

It is usually the case in real spreadsheet models that there are many unknown parameters, and it is often the case that there are multiple performance measures. This type of what-if analysis can become unwieldy when the number of parameters is more than a few. 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, maximum values. 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 increases to about 14 million, an infeasible amount of computing 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 is generally 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 using simulation for variation 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?

Generally, each cell in a spreadsheet model can be classified as containing one of three types of quantities:

- *Inputs to the model.* These can be parameters that are part of the model, such as unit costs or mean demand. They can also be the random variables that represent uncertain quantities in the model such as demand or price paid, or they can be the values of unknown parameters when one is doing a what-if analysis.
- *Intermediate computations.* These are the calculations in the model. For example, in an inventory model, these might compute the inventory levels or backlogs at the end of each period using the previous period's ending inventory and the sampled demand for the current period.
- *Outputs from the model.* These are 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. Some sample spreadsheet models are available at <http://seila.terry.uga.edu/spreadsheetSim>. It is difficult to describe all of the details in the process of setting up a spreadsheet simulation, but these models have annotations to help the reader to understand how they were constructed.

4.1 Setup the Model

The first step is to build the model in the spreadsheet using definite values for all parameters and uncertain quantities.

This step allows one to check the computations and assure the correctness of the model.

Second, replace the values in the cells that represent random or unknown quantities with formulas that sample their values from appropriate distributions. Appropriate formulas can be found in any reference on random variate generation. See Cheng (1998) for example. At this point, all random variates will be resampled each time the spreadsheet is recalculated.

Third, identify the "outputs" to the model. Actually, the modeler should know the desired performance measures when the model is created. Here, you want to identify the cells containing the values of these performance measures. At this point, you can observe the values of these cells change (i.e., be sampled) each time the spreadsheet is recalculated.

4.2 Create the Simulation Run

It is useful to distinguish two types of simulation experiments: (1) dynamic simulations that produce a time series of dependent observations and (2) static simulations that replicate the experiment independently, producing independent, identically distributed observations. The setup for each of these types is different.

In a dynamic simulation, the output values are observed periodically. 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 or intermediate computations. Often, the computations for each period can be placed in a single row in the spreadsheet. Then, the next period's values can be put in the next row and computed from the previous periods' cells. Once the computations for a representative set of periods are setup, 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 output observations in the simulation is contained in one or more columns of the spreadsheet.

When independent replications are performed, the model computations are usually contained in some region or regions of the spreadsheet. In order to collect multiple replications, we need to do a series of recalculations of the spreadsheet and save the outputs after each recalculation. The easiest way to do this is to create a Data Table and use it to tell the spreadsheet to go through an iterative recalculation, storing the values of the outputs between each recalculation. To setup the data table, select an area of the spreadsheet where the outputs will be stored and create 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, you place a formula that will contain the value of a specific output. Each row of this table will contain the replication

number and the outputs, i.e., observations on each performance measure, for that replication. Once this table outline has been created, you are ready to tell the spreadsheet to actually perform the replications.

Running the replications involves invoking a command such as the Data/Table command to tell the spreadsheet to perform the series of recalculations. This command was actually created to perform “what-if” scenarios by substituting each value in the first column into a specified cell, recalculating the spreadsheet, and recording the values of other cells that depend upon the substituted value. For example, you could use this technique to compute the present value of an investment as the rate of return varies from three percent to 10 percent. In our case, the contents of the cell in the first column of the data table is just a replication number, which is not actually used in the model calculations. However, the recalculation of the spreadsheet will cause all random variate sampling formulas to re-execute, producing new values for all random variates and thus new, statistically independent values for all outputs. The Data/Table command produces a dialog asking where you want to put the input value. You may select any unused cell and click "OK". The data table containing the results of the replications will fill quickly. When it is finished, each column of this table except the first will contain the set of observations on a specific performance measure. It is not difficult to do thousands, or even tens of thousands, of replications in this way.

4.3 Analyze the Data

Each simulation has its own analysis requirements (Alexopoulos and Seila 1998). For stochastic models, analysis normally involves applying statistical procedures to compute estimates of population parameters and confidence intervals for these estimates. When variational analysis is the objective, analysis is concerned with evaluating the *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.

Presentation generally includes some tables and graphs. Spreadsheets have extensive facilities that make it easy to produce these types of presentations in high quality.

5 SIMULATION ADD-INS FOR SPREADSHEETS

The process of developing and running a simulation in a spreadsheet can be simplified somewhat by us-

ing one of the available commercial add-in packages such as @RISK (<http://www.palisade.com>) or Crystal Ball (<http://www.decisioneering.com>). A free add-in package called SIMTOOLS.XLA by Professor Roger Myerson is also 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, ?). Thus, using the built-in RAND() function carries some risk that the random numbers will not behave as truly independent, random numbers. In @RISK and Crystal Ball, the random number generators have been tested and documented.

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, some random variates such as the Gamma and Weibull distributions are difficult or impossible to generate using 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.

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, (2) complex algorithms are difficult to implement, (3) spreadsheets are slower than some alternatives and (4) data storage is limited. Let's examine each of these limitations.

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. It is possible to build these structures in a spreadsheet, but they are contrived and inefficient.

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 is 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. So, models that require complex loops and other conditional computations may not work well in a spreadsheet.

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.

Finally, since the output data must be stored in the spreadsheet, usually in a column, the length of the output series is limited. In many spreadsheets, column lengths can be tens of thousands or even hundreds of thousands of cells. However, some models such as models 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 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

Spreadsheets provide a useful platform for many simulation models. The attractiveness of this platform comes from its availability, intuitive interface, ease of use and powerful features. At the same time, spreadsheets are not able to handle certain types of models that require complex data structures or extensive computing resources. Therefore, spreadsheets are not recommended for some simulation problems.

Many simulations do not need to be extensive. They are designed to provide ball-park estimates and to show general system behavior. This is true of many financial models. These models can usually be implemented most efficiently in a spreadsheet. Simulation problems for which spreadsheets are a useful platform also include relatively small prototype models which are designed to provide a proof of concept.

Commercial and free spreadsheets are continuing to be developed. Future versions will undoubtedly allow larger worksheets and be more efficient in the way that computations are performed. 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.
- Mattesich, R. 1961. Budgeting models and system simulation. *The Accounting Review* 36:384–397.

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 attending the Winter Simulation Conference since 1977 and served as Program Chair for the 1994 Orlando meeting. His professional interests include simulation methodology, especially output analysis, and applications in health care and finance. He is a member of ASA and INFORMS. His e-mail address is <andy@ms.terry.uga.edu>.