Proceedings of the 1993 Winter Simulation Conference
G.W. Evans, M. Mollaghasemi, E.C. Russell, W.E. Biles (eds.)

STATISTICAL ANALYSIS OF OUTPUT PROCESSES

John M. Charnes

Department of Management Science
University of Miami
Coral Gables, Florida 33124-8237

ABSTRACT

Discrete-event simulation output processes are seldom composed of independent and identically distributed random variables, but should be analyzed with statistical methods because they are stochastic. This tutorial presents the ideas underlying some of the methods that have been developed for statistical analysis of output processes. References to relevant work in the simulation literature and to standard textbooks on the subject are provided for those who are interested in deeper treatments of these methods than can be given in the space allotted here.

1 INTRODUCTION

Simulation is an approach to systems modeling in which the analyst hopes to obtain approximate answers to relevant questions about complex models rather than exact answers to questions about models with analytical solutions (such as those provided by queueing theory) that are rougher approximations to the system under study. Because the approach taken in simulation is to model the important components of a system as closely as possible using realistic structural assumptions and probability distributions for the important inputs, the outputs from a simulation model are generally sequences of random variables that have unknown probability distributions.

In one view, simulation can be seen as an attempt to estimate some parameters of the unknown output distributions, such as expected values, medians, or 95th percentiles. Thus, the analysis of simulation output is necessarily statistical, meaning that the simulation analyst must recognize the need to obtain accurate point estimates of the output measures of interest, as well as some measures of the precisions of these estimates. The precisions of estimates obtained from simulation models are usually given in the form of confidence intervals around the point estimates.

This tutorial is an overview of the ideas underlying some of the methods that have been developed for statistical analysis of simulation output processes. These processes usually require specialized techniques for analysis because they have characteristics that preclude the direct use of statistical methods that have been developed for independent and identically distributed (IID) random variables. It is assumed that the reader is familiar with basic statistical analysis of IID random variables, and is looking to this tutorial for an overview of what makes the analysis of simulation output different from basic statistical analysis.

Those interested in further treatment of the methods discussed here should consult one or more of the standard textbooks on discrete-event simulation, such as Banks and Carson (1984), Bratley, Fox, and Schrage (1987), Fishman (1978a, 1973), Kleijnen (1987, 1975, 1974), Law and Kelton (1991), or Lewis and Orav (1989); one or more of the survey works on simulation output analysis, such as Banks, Carson, and Goldsman (1990), Kleijnen (1982), Law (1983), Law and Kelton (1984, 1982), and Welch (1983); or the fine tutorials on simulation output analysis published in previous editions of the *Proceedings of the Winter Simulation Conference*, such as Goldsman (1992), or Seila (1990b).

2 TWO SYSTEM TYPES

There are two types of systems that call for different approaches for measuring the precisions of estimates obtained from the simulation models. In a *terminating* system the model has specific "start-up" and "shut-down" times. For example, a bank that opens at 9 A.M. and closes at 3 P.M. would typically be modeled as a terminating system. A *steady-state* system is one in which the model has no specific start-up or shut-down times. For example, a factory that operates twenty-four hours a day, seven days a week would typically be modeled as a steady-state system. In terminating systems, the simulation analyst is often interested in estimating parameters of the distributions of the outputs obtained for the period of time between start-up and shut-down, such as the mean time spent by customers inside the bank each day. In steady-state systems, the simulation analyst is often interested in estimating one or more steady-state parameters of the model. That is, the analyst assumes that if the model is in operation long enough, each output process will reach a state of statistical equilibrium, where the observations appear to fluctuate randomly according to some fixed but unknown distribution about which the analyst wishes to learn. If so, the output process will exhibit *stationarity*, meaning that certain characteristics of the underlying probability distribution (e.g., the mean, variance, and the fixed-lag correlations) will be invariant to the passage of simulated time. These two types of systems call for different methods of output analysis, as described below.

2.1 Terminating Systems

Suppose that the operation of a bank is to be simulated with a simple model where customers arrive to the bank, wait on line to be served if necessary, complete their transaction and then leave the bank. The analyst has set up the simulation model with the assumed interarrival time and service time probability distributions, and is now ready to run the model to estimate the mean time spent by customers (their "delay") inside the bank each day.

A good procedure begins with running the simulation model for a full day's operation, keeping track of each customer's delay so that D_1, the mean delay for the first day's customers, can be determined as the arithmetic average of the delays of all customers on the first day. This single value is the sole recorded output from the simulation of the first day's bank operation. The analyst then runs the simulation model for several more days, recording the mean delay for each day. After running the model for n days, the analyst obtains an output sequence $\{D_1, D_2, \ldots, D_n\}$. The elements of this output sequence are independent because independent pseudorandom numbers were used as inputs for each day's simulated operation. The elements of the output sequence are identically distributed because the model was unchanged from day to day; it is only because of the use of different pseudorandom numbers that different mean delays were obtained for each day. Thus, the output sequence $\{D_1, D_2, \ldots, D_n\}$ is amenable to analysis by statistical methods developed for IID random variables.

A point estimate for the daily mean customer delay in the bank is

$$\overline{D} = \sum_{i=1}^{n} D_i/n,$$

with the precision of this estimate given by the half-width of the $100(1-\alpha)\%$ confidence interval

$$\overline{D} \pm t_{1-\alpha/2}(n-1)\frac{\sqrt{\widehat{\mathrm{Var}}(D)}}{n},$$

where $t_{1-\alpha/2}(n-1)$ is the $100(1-\alpha/2)$th percentile of Student's t distribution with $n-1$ degrees of freedom, and

$$\widehat{\mathrm{Var}}(D) = \sum_{i=1}^{n} \frac{(D_i - \overline{D})^2}{n-1}.$$

This method, called the method of *independent replications*, is easily generalized. Instead of using the mean customer delay on day i as the output D_i in the ith replication of the model, the analyst could have recorded the 95th percentile each day, or the daily median customer delay (50th percentile). It is generally agreed that the method of independent replications is best for analyzing data obtained from terminating simulations because we can ensure that the output random variables are IID by using independent pseudorandom numbers as inputs for as many replications of an otherwise identical simulation model as we need to achieve any precision that is desired.

See Heidelberger and Lewis (1984), Iglehart (1976), Law (1980), and Seila (1982a, 1982b) for more information on terminating systems and percentile estimation.

2.2 Steady-State Systems

Now suppose that an analyst is simulating the factory mentioned above that operates twenty-four hours per day, seven days per week, and is interested in the time-weighted average number of parts in the system each day. Parts enter the factory, are transformed in some manner by the manufacturing process, and then leave the system. The simulation model is a representation of the manufacturing process.

The analyst codes the model to begin simulated time at midnight (denoted as time $t_0 = 0.0$) of the first day, at which time the initial number of parts in the system is counted and denoted as $N_1(t_0)$, where the subscript "1" on N indicates the first day of simulated time. Throughout the first day, every time a part enters or leaves the system the number of parts in the system is counted and recorded as $N_1(t_i)$, where t_i is the time of the ith event (an arrival to or departure from the system). Denote by e_1 the number of arrival and departure events occurring on day 1. In essence, what has been recorded is the continuous-time process for number in system on the first day ($N_1(t)$, for $0 \leq t \leq 24$), because this function changes only at times $t_1, t_2, \ldots, t_{e_1}$. The first element of the output sequence to be analyzed is found as

$$\begin{aligned} Y_1 &= \frac{1}{24}\int_0^{24} N_1(t)dt, \\ &= \sum_{j=1}^{e_1} N_1(t_{j-1})[t_j - t_{j-1}]/24 \\ &\quad + N_1(t_{e_1})[24 - t_{e_1}]/24, \end{aligned}$$

which is the time-weighted average number of parts in the system on the first day of simulated time.

The analyst then sets $N_2(0.0) = N_1(24.0)$, and records the number in system at each of the e_2 event times that the function $N_2(t)$ changes to obtain Y_2, the time-weighted average number of parts in system on the second day, in a manner analogous to that in which Y_1 was obtained. This is continued until the sequence of observations $\{Y_1, Y_2, \ldots, Y_n\}$ is obtained.

The elements of the sequence $\{Y_1, Y_2, \ldots, Y_n\}$ are not IID. Simulations of manufacturing systems typically need time to "warm up" before they reach steady-state, so the initial observations are unlikely to be distributed in a manner identical to the latter observations. This problem—*initialization bias*—is a consideration in virtually every simulation that requires steady-state output. The usual solution for the problem of initialization bias is to let the system warm up and then discard the initial observations, which are said to have come from the *initial transient* part of the simulation; the problem in practice is to decide how many initial observations to truncate in order to prevent the initial transient from affecting the output analysis adversely.

The problem of initialization bias and its potential remedies are discussed by Chance and Schruben (1992), Fishman (1972), Gafarian, Ancker, and Morisaku (1978), Heidelberger and Welch (1983), Kelton (1987, 1989), Kelton and Law (1983), Schruben (1982, 1981), Schruben, Singh, and Tierney (1983), and Wilson and Pritsker (1978a, 1978b).

Even if the random variables in the latter portion of the sequence $\{Y_1, Y_2, \ldots, Y_n\}$ are identically distributed, these variables are not likely to be independent, which is called the problem of *autocorrelation* in the data. The average number of parts in the system on one day is apt to be correlated with the average number of parts in the system on the previous day(s) because the system is not re-initialized each day. While the practice of not re-initializing is in general the most realistic thing to do for modeling purposes, the autocorrelation in the data imposes additional difficulties for output analysis. However, several remedies have been suggested in the simulation literature for the problem of autocorrelation in the data.

3 STEADY-STATE ANALYSIS

Much effort has been expended on research into the problem of how best to deal with the problem of autocorrelation in data obtained from simulations operating in steady-state. No single method has gained widespread acceptance within the simulation community for use on all models, and research into these methods is ongoing. Some of the better known methods are described briefly below.

3.1 Independent Replications With Truncation

This method is similar to the method of independent replications for terminating systems. In the *independent replications with truncation* method, the steady-state system is run several times starting from the same initial conditions, but using independent pseudorandom numbers for each replication. Usually the analyst makes pilot runs and uses one of the methods mentioned above to determine approximately when the initial transient effect can be ignored, truncates the output up to that point, and then uses the remaining elements in the output sequence for analysis. This method has been viewed by some in the past as too wasteful of computer time, but it is becoming more viable as the cost of computing continues to plummet. The advantage of this technique is that independent observations are sure to be obtained; the disadvantage is the "wasting" of more observations than does any of the alternative *one-long-run* methods, for which the analyst generates a large number of observations with only one initialization of the system. Using a one-long-run method makes it necessary to truncate only one set of initial values to mitigate the effects of intitialization bias, but then forces the analyst to deal with the problem of autocorrelation in the remainder of the output sequence. See Whitt (1991, 1989) for further discussion of these issues.

3.2 Batch Means Method

The batch means method is a one-long-run attempt to deal with autocorrelation in the data by combining adjacent autocorrelated observations in the output sequence into (nearly) uncorrelated batches. The mean of each batch is calculated and the collection of these means is then treated as a sequence of IID observations. Assume that the observations generated during the initial transient period have been truncated and the observations have been renumbered so that the output sequence $\{Y_1, Y_2, \ldots, Y_n\}$ is observed while the simulation model is operating in steady-state. The sequence is divided into m batches of k consecutive observations, where $n = mk$. The mean of batch i is

$$X_i = \frac{1}{k} \sum_{j=1+(i-1)k}^{ik} Y_j$$

for $i = 1, \ldots, m$ and the overall sample mean of the output sequence is

$$\overline{X} = \frac{1}{m} \sum_{i=1}^{m} X_i = \frac{1}{n} \sum_{j=1}^{n} Y_j,$$

while a $100(1-\alpha)\%$ confidence interval for the mean is

$$\overline{X} \pm t_{1-\alpha/2}(m-1) \frac{\sqrt{\widehat{\mathrm{Var}}(X)}}{m},$$

where $t_{1-\alpha/2}(m-1)$ is the $100(1-\alpha/2)$th percentile of Student's t distribution with $m-1$ degrees of freedom, and

$$\widehat{\mathrm{Var}}(X) = \sum_{i=1}^{m} \frac{(X_i - \overline{X})^2}{m-1}.$$

See Fishman (1978a), Law (1977), Schmeiser (1982), and Schmeiser and Kang (1981) for more information on the batch means method.

3.3 Advanced Steady-State Techniques

Several other methods have been proposed for analysis of steady-state output processes. Proper use of these methods relies on the validity of different assumptions than made so far here about the output processes, and requires a higher degree of statistical sophistication on the part of the reader than has been assumed here. These methods are not recommended for novice analysts, but are mentioned briefly here for informational purposes.

Regenerative Method. The regenerative method calls for identifying random times at which the output process regenerates (begins anew probabilistically). The observations between regeneration points can be used to obtain independent random variables, which can then be analyzed in a manner similar to one of those given above. The interested reader should consult Crane and Iglehart (1975), Crane and Lemoine (1977), Fishman (1977), and Iglehart (1978).

Spectral Analysis Method. The spectral analysis method of estimating the variance of the sample mean of an output process relies on the estimation of the spectral density function of the process. The spectral density function of a process is the Fourier transform of the autocorrelation function of the process, and can be used to form a confidence interval on the mean of the process. This frequency-domain, time-series technique is discussed further in Duket and Pritsker (1978), and Heidelberger and Welch (1981a, 1981b).

Standardized Time Series. The standardized time series method generalizes the notion of standardizing normal random variables to standardizing output processes. The standardized process is then used to form a confidence interval for the mean of the process. See Glynn and Iglehart (1990), Goldsman and Schruben (1990), and Schruben (1983) for more information.

Autoregressive Method. This method involves fitting an autoregressive model to the output sequence and using the estimated parameters of the model to estimate the variance of the sample mean. This variance estimate is then used to form a confidence interval on the mean. This time-domain, time-series technique is discussed further in Fishman (1978b), and Schriber and Andrews (1984).

4 SEQUENTIAL METHODS

The methods described above are called *fixed-sample-size methods* because the simulation is run until the specified number of observations n is generated, at which time a confidence interval is obtained that will have some *random amount of precision* that is determined by the half-width. In *sequential methods* the precision is specified in advance by placing an upper bound on the half-width and running the simulation to obtain whatever *random number of observations* is necessary to meet or exceed the specified precision. The use of sequential methods places additional demands on the simulation coding; see Fishman (1977), Heidelberger and Welch (1983), Law and Kelton (1982), and Law and Carson (1979) for more information.

5 COMPARING SYSTEMS

Arguably, it can be said that the greatest benefits of simulation come from the ability to compare models of alternative systems before deciding which actual system to put in place. The literature on comparing systems can be divided roughly into two topics: comparing two systems, and selecting the best subset (possibly a singleton) from a group of several competing systems. See Chapter 10 of Law and Kelton (1991) and the references therein for an excellent introduction to these topics.

6 VARIANCE REDUCTION

Simulation experiments are unique in that the source of randomness can be controlled by the experimenter. The pseudorandom numbers that drive the models can sometimes be manipulated in ways that help to increase the precision of the estimates obtained from the model by reducing the variance of the estimates, while still maintaining the random characteristics necessary for a meaningful experiment. These *variance reduction* methods can be a powerful method of gaining more information from a simulation model.

The variance reduction methods using common random numbers and antithetic variates rely on the relationship

$$\mathrm{Var}(a\overline{X} + b\overline{Y}) = a^2\mathrm{Var}(\overline{X}) + b^2\mathrm{Var}(\overline{Y}) + 2ab\mathrm{Cov}(\overline{X}, \overline{Y}), \quad (1)$$

where

$$\mathrm{Var}(Z) = \mathrm{E}[(Z - \mathrm{E}[Z])^2],$$

$$\mathrm{Cov}(U, V) = \mathrm{E}[(U - \mathrm{E}[U])(V - \mathrm{E}[V])],$$

and $\mathrm{E}[\cdot]$ is the expectation operator.

Common random numbers are typically used for comparing two or more systems and depend on the induction of positive correlation between estimates $\overline{X}$ and $\overline{Y}$ obtained from different systems. The simplest use of *antithetic variates* depends on the induction of negative correlation between estimates $\overline{X}$ and $\overline{Y}$ obtained from different simulations of the same system.

6.1 Common Random Numbers

Using common random numbers can be as simple as running two alternative systems with the same pseudorandom inputs in the hope that observed differences in performances are due to differences in the systems, and not due to the randomness of the inputs. The idea is to compare alternatives under the same experimental conditions. In comparing two factory layouts, it seems reasonable to think that a good way to compare them is to compare the average flow (sojourn) times for the same jobs going through each layout.

In a common random numbers scheme, two layouts are simulated with the same n jobs arriving at the same times to each layout, with the flow times for jobs going through the first layout recorded as $\{X_1, X_2, \ldots, X_n\}$; and the flow times for jobs going through the second layout recorded as $\{Y_1, Y_2, \ldots, Y_n\}$. An estimate of the difference between average flow times for the two layouts is

$$\overline{X} - \overline{Y} = \sum_{i=1}^{n}(X_i - Y_i)/n.$$

The precision of this point estimate depends on $\mathrm{Var}(\overline{X} - \overline{Y})$, which is given by (1) with $a = 1$ and $b = -1$.

The use of common random numbers to generate the same arrivals to each layout implies that $\mathrm{Cov}(\overline{X}, \overline{Y}) > 0$, which means that $\mathrm{Var}(\overline{X} - \overline{Y})$ is smaller with common random numbers than it would be if independent pseudorandom numbers were used to generate independent arrivals to each layout (in which case $\mathrm{Cov}(\overline{X}, \overline{Y}) = 0$). Thus the use of common random numbers gives a more precise point estimate. However, note that if the analyst was not careful in *synchronizing* the models—using the same pseudorandom numbers for the same purpose in each layout—and inadvertently coded the model so that $\mathrm{Cov}(\overline{X}, \overline{Y}) < 0$, the use of common random numbers would give a less precise point estimate than would the use of independent pseudorandom numbers.

6.2 Antithetic Variates

The antithetic variates method of variance reduction is used to increase the precision of estimates obtained from a single run of a simulation model. The idea is to try to obtain pairs of observations on a single measure from two different runs of a simulation model in which an observation that is above the true parameter value in one run is countered with an estimate below the true parameter value in the second run. Then the sequence of averages of the pairs of estimates will have smaller variance than either the sequence of the first observations or the sequence of the second observations.

In a simple antithetic variates scheme, the sequence of flow times for jobs going through a factory layout, $\{X_1, X_2, \ldots, X_n\}$, is obtained by using pseudorandom numbers $\{U_1, U_2, \ldots, U_m\}$ to generate random variates that are used for different purposes in the model (such as interarrival and service times). Then the same simulation model is run to obtain flow times $\{Y_1, Y_2, \ldots, Y_n\}$ by using the complementary pseudorandom numbers $\{1 - U_1, 1 - U_2, \ldots, 1 - U_m\}$ in the same places that $\{U_1, U_2, \ldots, U_m\}$ were used in the first run (i.e., the models must be synchronized). An estimate of the average flow is then obtained as

$$\frac{\overline{X}+\overline{Y}}{2}=\sum_{i=1}^{n}(X_i+Y_i)/(2n).$$

The precision of this point estimate depends on $\mathrm{Var}(\frac{\overline{X}+\overline{Y}}{2})$, which is given by (1) with $a=\frac{1}{2}$, and $b=\frac{1}{2}$.

The use of complementary pseudorandom numbers to generate the pairs of observations on the layout implies that $\mathrm{Cov}(\overline{X},\overline{Y})<0$, which means that $\mathrm{Var}(\frac{\overline{X}+\overline{Y}}{2})$ is smaller with the antithetic variates scheme than it would be if independent pseudorandom numbers were used to generate pairs of observations on the layout (in which case again $\mathrm{Cov}(\overline{X},\overline{Y})=0$). Thus the use of antithetic variates gives a more precise point estimate. Note that here, if the analyst didn't synchronize the models, and inadvertently coded the model so that $\mathrm{Cov}(\overline{X},\overline{Y})>0$, the use of antithetic variates would give a less precise point estimate than would the use of independent pseudorandom numbers.

6.3 Other Variance Reduction Methods

Researchers are continuing to investigate the better use of both the common random numbers and antithetic variates methods of variance reduction as well as their combined usage in a single experiment. Other variance reduction methods include *control variates*—taking advantage of correlation between estimates and other selected random variables to reduce estimator variance, *indirect estimation*—using known theoretical relations among estimates to reduce estimator variance, and *conditioning*—exploiting some special property of a model to replace an estimate by a known parameter value. Interested readers should begin with Chapter 11 of Law and Kelton (1991), Nelson (1992, 1987), and Wilson (1984).

7 MULTIVARIATE PROCESSES

The methods described thus far have been concerned with analyzing single measures of performance, or *univariate* output processes, but simulation analysts will often be concerned with multiple measures of performance, or *multivariate* output processes. For example, an analyst might be interested in estimating the mean customer delay in a bank for each hour of the day (e.g., 9:00–9:59AM, 10:00–10:59AM, etc.) or estimating the mean number of jobs at each of several workstations in a factory. However, these measures are almost always correlated.

Analyzing multivariate output processes sometimes requires that the correlation among the performance measures be taken into consideration. The precision for these correlated performance measures can be given in the form of multi-dimensional confidence regions. Many of the univariate methods given above have been generalized to the multivariate case; see Charnes (1990), Charnes and Kelton (1993, 1988), Chen and Seila (1987), Seila (1990a, 1990b, 1984), and Yang and Nelson (1988) for details.

A practical problem with using confidence regions that account for correlation among performance measures is that the regions are hard to interpret. The Bonferroni Inequality provides the basis for one straightforward method of constructing several univariate confidence intervals simultaneously. See Charnes and Kelton (1993, 1988) for comparisons of simultaneous confidence intervals and other confidence regions.

8 CONCLUSION

Simulation experiments must be replicated several times before the analyst should be willing to believe that the obtained results are meaningful. Because simulation output processes are random, statistical techniques should be used for analysis of the output data. However, "raw" simulation output processes are seldom (if ever) IID.

Most of the methods of analysis discussed here involve transforming the raw output processes in some manner so that statistical techniques developed for IID random variables can be used to analyze the transformed processes. These methods are designed to yield point estimates of performance measures along with a measure of precision of the point estimate. The measure of precision is usually given in the form of a confidence interval.

This tutorial is intended as a brief overview of the ideas underlying some of the specialized techniques that have been developed for analyzing simulation output processes. Those readers wishing to learn more about the details of these techniques should start with one of the standard textbooks listed in §1. Those interested in digging deeper into the methods should peruse the references given here for each method, as well as the references provided in the standard textbooks.

ACKNOWLEDGEMENT

The author thanks the University of Miami for providing support in the form of a James W. McLamore Summer Award in Business and the Social Sciences, the Minnesota Supercomputer Institute for providing support and the equipment with which to prepare this paper, and the Department of Operations and Management Science at the University of Minnesota for their gracious hospitality.

REFERENCES

Banks, J., and J. S. Carson, II. 1984. *Discrete-Event System Simulation.* Englewood Cliffs, New Jersey: Prentice-Hall.

Banks, J., J. S. Carson, II, and D. Goldsman. 1990. Computer simulation. In *Handbook of statistical methods for engineers and scientists*, ed. H. M. Wadsworth, 12.1–12. 36. New York: McGraw-Hill.

Bratley, P., B. L. Fox, and L. E. Schrage. 1987. *A Guide to Simulation*, Second Edition. New York: Springer-Verlag.

Chance, F., and L. W. Schruben. 1992. Establishing a truncation point in simulation output. Technical Report, School of Operations Research and Industrial Engineering, Cornell University, Ithaca, New York.

Charnes, J. M. 1990. Power comparisons for the multivariate batch-means method, In *Proceedings of the 1990 Winter Simulation Conference*, eds. O. Balci, R. P. Sadowski, R. E. Nance, 281–287. Institute of Electrical and Electronics Engineers, New Orleans, Louisiana.

Charnes, J. M., and W. D. Kelton. 1993. Multivariate autoregressive techniques for constructing confidence regions on the mean vector, *Management Science*, in press.

Charnes, J. M., and W. D. Kelton. 1988. A comparison of confidence region estimators for multivariate simulation output, In *Proceedings of the 1988 Winter Simulation Conference*, eds. M. Abrams, P. Haigh, J. C. Comfort, 458–465. Institute of Electrical and Electronics Engineers, San Diego, California.

Chen, D. R. and A. F. Seila. 1987. Multivariate inference in stationary simulation using batch means, In *Proceedings of the 1987 Winter Simulation Conference*, eds. A. Thesen, H. Grant, W. D. Kelton, 302–304. Institute of Electrical and Electronics Engineers, Atlanta, Georgia.

Crane, M. A., and D. L. Iglehart. 1975. Simulating stable stochastic systems: III. Regenerative processes and discrete–event simulations. *Operations Research* 23:33–45.

Crane, M. A. and A. J. Lemoine. 1977. *An Introduction to the Regenerative Method for Simulation Analysis*, Springer-Verlag, New York, NY.

Duket, S. D. and A. A. B. Pritsker. 1978. Examination of Simulation Output Using Spectral Methods. *Math. Comput. Simulation* 20: 53–60.

Fishman, G. S. 1978a. Grouping Observations in Digital Simulation. *Management Science* 24: 510–521.

Fishman, G. S. 1978b. *Principles of Discrete Event Simulation*, John Wiley, New York, NY.

Fishman, G. S. 1977. Achieving Specific Accuracy in Simulation Output Analysis. *Communications of the Association for Computing Machinery* 20: 310-315.

Fishman, G. S. 1973. *Concepts and Methods in Discrete Event Digital Simulation*, John Wiley, New York, NY.

Fishman, G. S. 1972. Bias Considerations in Simulation Experiments. *Operations Research* 20: 785–790.

Gafarian, A. V., C. J. Ancker, Jr., and T. Morisaku. 1978. Evaluation of commonly used rules for detecting 'steady state' in computer simulation. *Naval Research Logistics Quarterly* 25:511–529.

Glynn, P. W., and D. L. Iglehart. 1990. Simulation output analysis using standardized time series. *Mathematics of Operations Research* 15:1–16.

Goldsman, D. 1992. Simulation output analysis. In *Proceedings of the 1992 Winter Simulation Conference*, eds. J. J. Swain, D. Goldsman, R. C. Crain, J. R. Wilson. Institute of Electrical and Electronic Engineers, Arlington, Virginia, 97–103.

Heidelberger, P. and P. A. W. Lewis. 1984. Quantile Estimation in Dependent Sequences. *Operations Research* 32: 185–209.

Heidelberger, P., and P. D. Welch. 1983. Simulation run length control in the presence of an initial transient. *Operations Research* 31:1109–1144.

Heidelberger, P., and P. D. Welch. 1981a. A spectral method for confidence interval generation and run length control in simulations. *Communications of the ACM* 24:233–245.

Heidelberger, P. and P. D. Welch. 1981b. Adaptive Spectral Methods for Simulation Output Analysis. *IBM Journal of Research and Development* 25: 860–876.

Iglehart, D. L. 1978. The Regenerative Method for Simulation Analysis. In *Current Trends in Programming Methodology, Vol. III. Software Engineering*, K. M. Chandy and R. T. Yeh, Eds. Prentice-Hall, Englewood Cliffs, NJ, 52–71.

Iglehart, D. L. 1976. Simulating Stable Stochastic Systems: VI. Quantile Estimation. *Journal of the Association for Computing Machinery* 23: 347–360.

Kelton, W. D. 1989. Random Initialization Methods in Simulation. *IIE Transactions* 21: 355–367.

Kelton, W. D. 1987. Stochastic Simulation: Initial Transient Techniques. *Systems and Control Encyclopedia; Theory, Technology, Applications*, M. G. Singh, Ed. Pergamon Press, Oxford, 4651–4654.

Kelton, W. D. and A. M. Law. 1983. A New Approach for Dealing with the Startup Problem in Discrete Event Simulation. *Naval Research Logistics Quarterly* 30: 641–658.

Kleijnen, J. P. C. 1987. *Statistical Tools for Simulation Practitioners.* Marcel Dekker, Inc., New York.

Kleijnen, J. P. C. 1982. Statistical Aspects of Simulation: An Updated Survey. *Statistica Neerlandica* 36: 165–186.

Kleijnen, J. P. C. 1975. *Statistical Techniques in Simulation, Part II*, Marcel Dekker, New York.

Kleijnen, J. P. C. 1974. *Statistical Techniques in Simulation, Part I*, Marcel Dekker, New York.

Law, A. M. and W. D. Kelton. 1982. Confidence Intervals for Steady-State Simulations, II: A Survey of Sequential Procedures. *Management Science* 28: 550–562.

Law, A. M. 1980. Statistical Analysis of the Output Data from Terminating Simulations. *Naval Research Logistics Quarterly* 27: 131–143.

Law, A. M. 1983. Statistical analysis of simulation output data. *Operations Research* 31: 983–1029.

Law, A. M. 1977. Confidence Intervals in Discrete Event Simulation: A Comparison of Replication and Batch Means. *Naval Research Logistics Quarterly* 24: 667–678.

Law, A. M. and J. S. Carson. 1979. A Sequential Procedure for Determining the Length of a Steady-State Simulation. *Operations Research* 27: 1011–1025.

Law, A. M. and W. D. Kelton. 1991. *Simulation Modelling and Analysis*, Second Edition. McGraw-Hill: New York.

Law, A. M. and W. D. Kelton. 1984. Confidence Intervals for Steady-State Simulations: I. A Survey of Fixed Sample Size Procedures. *Operations Research* 32: 1221–1239.

Lewis, P. A. W. and E. J. Orav. 1989. *Simulation Methodology for Statisticians, Operations Analysts, and Engineers, Volume I*, Wadsworth, Inc., Belmont, CA.

Nelson, B. L. 1992. Designing efficient simulation experiments. In *Proceedings of the 1992 Winter Simulation Conference*, eds. J. J. Swain, D. Goldsman, R. C. Crain, J. R. Wilson. Institute of Electrical and Electronic Engineers, Arlington, Virginia, 126–132.

Nelson, B. L. 1987. Variance reduction for simulation practitioners. In: *Proceedings of the 1987 Winter Simulation Conference* (A. Thesen, H. Grant, W. D. Kelton, eds.) Institute of Electrical and Electronic Engineers, San Francisco, California, 43–51.

Schmeiser, B. W. 1982. Batch size effects in the analysis of simulation output. *Operations Research* 30:556–568.

Schmeiser, B. W. and K. Kang. 1981. Properties of Batch Means from Stationary ARMA (1,1) Time Series. Technical Report 81-3, School of Industrial Engineering, Purdue University, West Lafayette, IN.

Schriber, T. J. and R. W. Andrews. 1984. ARMA-Based Confidence Intervals for Simulation Output Analysis. *American Journal of Mathematical and Management Sciences* 4: 345–374.

Schruben, L. W. 1983. Confidence interval estimation using standardized time series. *Operations Research* 31:1090–1108.

Schruben, L. W. 1982. Detecting initialization bias in simulation output. *Operations Research* 30:569–590.

Schruben, L. W. 1981. Control of initialization bias in multivariate simulation response. *Communications of the ACM*, 24(4): 246–252.

Schruben, L. W., H. Singh, and L. Tierney. 1983. Optimal tests for initialization bias in simulation output. *Operations Research* 31: 1167–1178.

Seila, A. F. 1990a. Multivariate Estimation of Conditional Performance Measures in Regenerative Simulation. *American Journal of Mathematical and Management Sciences* 10: 17–50.

Seila, A. F. 1990b. Output Analysis for Simulation. In *Proceedings of the 1990 Winter Simulation Conference*, eds. O. Balci, R. P. Sadowski and R. E. Nance, 49–54. The Society for Computer Simulation, International, San Diego, CA.

Seila, A. F. 1984. Multivariate Simulation Output Analysis. *American Journal of Mathematical and Management Sciences* 4: 313–334.

Seila, A. F. 1982a. A Batching Approach to Quantile Estimation in Regenerative Simulations. *Management Science* 28: 573–581.

Seila, A. F. 1982b. Percentile Estimation in Discrete-Event Simulation. *Simulation* 39: 193–200.

Welch, P. D. 1983. The Statistical Analysis of Simulation Results. *The Computer Performance Modeling Handbook*, S. S. Lavenberg, Ed., Academic Press, New York, 268–328.

Whitt, W. 1991. The efficiency of one long run versus independent replications in steady-state simulation. *Management Science* 37: 645–666.

Whitt, W. 1989. Planning queueing simulations. *Management Science* 35: 1341–1366.

Wilson, J. R. 1984. Variance Reduction Techniques for Digital Simulation. *American Journal of Mathematical and Management Sciences* 4: 277–312.

Wilson, J. R. and A. A. B. Pritsker. 1978a. A Survey of Research on the Simulation Startup Problem. *Simulation* 31: 55–58.

Wilson, J. R. and A. A. B. Pritsker. 1978b. Evaluation of Startup Policies in Simulation Experiments. *Simulation* 31: 79–89.

Yang, W.-N. and B. L. Nelson. 1988. Multivariate estimation and variance reduction in terminating and steady-state simulation, In *Proceedings of the 1988 Winter Simulation Conference*, eds. M. Abrams, P. Haigh, J. C. Comfort, 466–472. Institute of Electrical and Electronics Engineers, San Diego, California.

AUTHOR BIOGRAPHY

JOHN M. CHARNES is Assistant Professor in the Management Science Department at the University of Miami in Coral Gables, Florida. His main research interests are in multivariate time series and simulation output analysis. He is Associate Editor for the *TIMS College on Simulation Newsletter*.