

ALTERNATIVE APPROACHES FOR SPECIFYING INPUT DISTRIBUTIONS AND PROCESSES
(PANEL)

Chair

W. David Kelton

Department of Operations and Management Science
Curtis L. Carlson School of Management
University of Minnesota
Minneapolis, Minnesota 55455

Panelists

Bennett L. Fox

Department of Mathematics
University of Colorado – Denver
Campus Box 170, 1200 Larimer Street
Denver, Colorado 80204-5300

Mark E. Johnson

School of Industrial and Systems
Engineering
Georgia Institute of Technology
Atlanta, Georgia 30332-0205

Averill M. Law

Simulation Modeling and Analysis
Company
P.O. Box 40996
Tucson, Arizona 85717

Bruce W. Schmeiser

School of Industrial Engineering
Purdue University
West Lafayette, Indiana 47907

James R. Wilson

School of Industrial Engineering
Purdue University
West Lafayette, Indiana 47907

Discussants

John Meszaros

Systems Modeling Corporation
The Park Building
504 Beaver Street
Sewickley, Pennsylvania 15143

Cynthia L. Morey

C-P-C Simulation
General Motors of Canada Limited
MC:052-001, 1908 Colonel Sam Drive
Oshawa, Ontario L1H 8P7, Canada

Susan E. Romens

3M Plant Engineering
Building 42-3W-04
900 Bush Avenue, P.O. Box 33331
St. Paul, Minnesota 55133-3331

ABSTRACT

One way to think about the process of building a simulation model is to classify the modeling into two (probably overlapping) activities. *Structural* modeling involves setting up the basic physical and logical structure, without bothering with specifics of parameter values, input-distribution forms, or run lengths, etc. For example, in a model of a communications system we would have to specify the types of messages, network topology, and protocols. In *quantitative* modeling we are forced to get specific about parameter values, forms of input-probability distributions, as well as the numerical values of the parameters of those distributions. In the communications example, we would have to specify distributions for the lengths of the different types of messages and a stochastic process for the frequency of request of network services from its nodes.

Typically, most attention is focused on structural modeling. Certainly, it is important to get that part right. Also, of the two, it may in some ways be more fun.

However, accurate quantitative modeling is just as essential. Examples abound to illustrate the grave errors that can result from using an "incorrect" probability distribution, even one that might be "correct" for the expected value and even variance. Furthermore, in large simulations there could well be hundreds of different input distributions to specify, each representing a separate step in a fabrication process. Practitioners have consistently identified the problem of input-distribution specification as being among the most difficult aspects of a simulation study.

Traditionally, simulation analysts have used well-known statistical techniques to "fit" a "standard" distribution to observed data; in terms of generating variates from such distributions, most of the simulation languages provide adequate support. For instance, an exponential, gamma, or Weibull distribution could be fitted to observed service-time data. However, there has been debate over the degree to which this is appropriate. As a diametrically opposed view, it has

been proposed that empirical distributions (specified directly from the data, and perhaps modified in the tails) be used instead. Positions somewhere in between these two extremes have also been proposed, using flexible "families" of distributions with a fairly large number of parameters (4 or 5) to allow a wide range of shapes and ranges.

This panel discussion will provide a forum for discussion of philosophical issues, but more importantly will shed light on their quantitative relative merits. Researchers in the area will give position statements and evidence, and discussants from the simulation-practitioner community will comment on the alternative approaches, including their practicality.

This paper presents (in alphabetical order) brief position statements written by the panelists to serve as a point of departure for the discussion.

BENNETT L. FOX

This position statement on input-distribution modeling is an update of the views expressed in Bratley et al. [1987, chapter 4].

Keep the model simple but not simplistic. If there are lot of parameters, estimation may require intricate numerical routines and may well be statistically unreliable—especially when the sample contains observations far out in the tail. If there are too few parameters, getting a good fit may be impossible.

Choose the model so that parameter estimation is easy and variate generation is fast, but not at the price of using a simplistic model. If variate generation by inversion is fast, variance reduction by correlation induction is greatly facilitated.

Unless some distribution is extremely peaked or is supported on an interval which can't be taken, to a first approximation, as zero to infinity, use phase-type distributions to reduce the model to a continuous-time Markov chain. Variance reduction techniques and gradient

estimation are easier in that setting; e.g., see Fox and Glynn [1990] and Glasserman [1990], respectively. In addition, transitions can be generated quickly; see Fox [1990] and Fox and Young [1990]. Also, the (useful) simterpolation technique of Reiman et al. [1990] requires at least one phase-type distribution—though not necessarily a complete Markov-chain model. Since I have not worked on estimating phase-type distributions, I'll leave a discussion of that to someone who has; my feeling is that there is a lot to say on that subject that has not yet been said. To check some of what has been said, see several papers cited in Altiok [1989].

If fitting a continuous-time Markov chain model as above doesn't work, use a linearly-interpolated empirical distribution—possibly with a (truncated) exponential tail as in Bratley, Fox, and Schrage [1987]. "Coalesce" the breakpoints by letting the (subsequently-interpolated) "empirical" distribution jump only at the k th, $2k$ th, ... order statistics where k is chosen as a function of the number n of observations to make the number of jumps at most a few dozen. Devroye [1986, page 767] suggests choosing k as a function of n to get a consistent density estimate, of interest when n is large and (more importantly) when computer memory is virtually unconstrained.

Nonstationarity can be hard to model. If only one distribution is non-stationary and a continuous-time Markov-chain model is appropriate, a piecewise-constant intensity corresponding to that distribution may be suitable. Given the breakpoints, the intensity function is easy to estimate. Fox and Glynn [1990, section 6] show how to handle the resulting process efficiently.

Modeling multivariate distributions is also tough. In small dimensions, I'd look at generalizing the quasi-empirical approach in Bratley, Fox, Schrage [1987]; beyond three or four dimensions, however, I'd guess that this is a dead end. In higher dimensions, I'd look at families of distributions indexed by a modest number of parameters; Johnson [1987] surveys these nicely. These families offer significant modeling flexibility, but in small dimensions—certainly in one dimension—they fail the simplicity test relative to the quasi-empirical approach. With the latter, parameter estimation is easier and variate generation is faster.

REFERENCES

- Altiok, T. (1989), "Approximate Analysis of Queues in Series with Phase-Type Service Times and Blocking," *Operations Research* 37, 601–610.
- Bratley, P., B.L. Fox, and L.E. Schrage (1987), *A Guide to Simulation*, Springer-Verlag, New York, NY.
- Devroye, L. (1986), *Non-Uniform Random Variate Generation*, Springer-Verlag, New York, NY.
- Fox, B.L. (1990), "Generating Markov-Chain Transitions Quickly: I," *ORSA Journal on Computing* 2, 126–135.
- Fox, B.L. and P.W. Glynn (1990), "Discrete-Time Conversion for Simulating Finite-Horizon Markov Processes," *SIAM Journal on Applied Mathematics*, to appear.
- Fox, B.L. and A.R. Young (1990), "Generating Markov-Chain Transitions Quickly: II," technical report.
- Glasserman, P. (1990), "Derivative Estimates from Simulation of Continuous-Time Markov Chains," Technical Report, AT&T Bell Laboratories, Holmdel, NJ.
- Johnson, M.E. (1987), *Multivariate Statistical Simulation*, Wiley, New York, NY.
- Reiman, M.I., B. Simon, and J.S. Willie (1990), "Simterpolation: A Simulation-Based Interpolation Approximation for Queuing Systems," Technical Report, AT&T Bell Laboratories, Murray Hill, NJ.

MARK E. JOHNSON

I am delighted to be given the opportunity to provide some random thoughts on the topic "alternative approaches for specifying input distributions and processes." I will proceed by reacting to the thought piece provided by our distinguished organizer.

The debate between using smooth data summaries embodied by well-known probability distributions (fit to the data) versus empirical cdf's with tail adjustments parallels the debates I heard recently at a

bootstrapping conference. That is, should we be using the bootstrap as it was originally interpreted (re-sampling) or should we use a parametric bootstrap. The bootstrapper crowd seems to put substantial weight of evidence on asymptotics. Fortunately, the same criterion has not yet strongly infiltrated the simulation crowd (although there is always the threat). So much for a sound bite perspective. A rather more pertinent point seems to be always overlooked regardless of the crowd. It hardly matters what distribution estimator one uses if the associated data are irrelevant. To illustrate, consider the proverbial bank-teller problem. The observed bank-teller data for a system in need of considerable improvement could be nearly useless. Slow tellers lead to few arrivals so that even an exact representation of this system and distribution fits could be deceptive to future system configurations. One would hope that arrival rates would pick up if the service in turn improves. A problem with this is that the current situation does not provide the data to which analysts are eagerly aspiring to subject their tools.

I have never cared for the (simplistic) view that data get collected, fit using some standard distribution type, and then blessed or blasted by means of a goodness-of-fit test. If blessed, the reason could be bad (too small a data set or as above, the data set is the 'wrong' one). If blasted, the practical lack of fit may be so small as to be not worth bothering with. The philosophical point to be made is that we want to be able to tell how far from the truth we might be rather than supplying a simple yes-no response for a model.

I would argue that it is not difficult to specify distributions in simulation studies. In fact, it is really a rather straightforward operation. The problem is defending those that have been selected. Arguments such as 'they passed some hypothesis test' or that 'many arrival patterns are exponential' are not terribly compelling. Perhaps it is a mistake to look for the one fixed static once-and-for-all set of distributions but rather consider the problem to be dynamic. Once the bank model has stumbled to some improved performance, further improvements should be sought. Continuous quality improvement is the rallying cry of the 90s, after all. Perhaps attention should be directed to using the input distributions to assess how much the current system can take and when further small/large refinements are called for. A missing ingredient in many simulation studies seems to be the role of the manager/decision maker. I hope we aren't just simulating systems to see if the software can mimic reality.

AVERILL M. LAW

There are two major situations experienced by the simulation practitioner: (1) System data are available on the input random variable of interest (e.g., machine repair times in a factory); (2) No system data exist on the random variable. We now discuss each of these cases in some detail.

Suppose that system data X_1, X_2, \dots, X_n are available and we would like to specify an "appropriate" corresponding probability distribution to use in a simulation model. If it is possible to "fit" a standard (theoretical) probability distribution (e.g., exponential, gamma, or lognormal) to the data that provides a "good" representation, then we believe that this is generally the best modeling approach. (See Law and Kelton [1991, pp. 356-400] for a three-activity approach for specifying a theoretical distribution.) If no standard distribution works well, then we recommend using an empirical distribution based on the X_i 's (see Law and Kelton [1991, pp. 350-353]). An important practical advantage of using standard and empirical distributions is that they are supported by most simulation software.

The following are some reasons why we believe that a good theoretical distribution is generally preferable to an empirical distribution:

- An empirical distribution may have certain "irregularities," particularly if only a small number of data values is available. A theoretical distribution, on the other hand, "smooths out" the data and may provide information on the overall underlying distribution.
- If empirical distributions are used in the usual way, it is not possible to generate values outside the range of the observed data in the simulation. This is unfortunate, since many measures of performance for simulated systems depend heavily on the

probability of an "extreme" event's occurring, e.g., generation of a very large service time. With a fitted theoretical distribution, on the other hand, values outside the range of the observed data can be generated.

There may be a compelling physical reason in some situations for using a certain theoretical distribution form as a model for a particular input random variable. For example, interarrival times of customers to a service facility are often exponentially distributed if arrivals occur at a constant rate.

A theoretical distribution is a compact way of representing a set of data values. Conversely, if n data values are available from a continuous distribution, then $2n$ values (data and corresponding cumulative probabilities) must be entered and stored in the computer to represent an empirical distribution in many simulation languages. Thus, use of an empirical distribution will be cumbersome if the data set is large.

Note that in many cases we do not expect any theoretical distribution to be an *exact* representation of the true underlying random variable of interest. Instead, we are trying to find a theoretical distribution that is good enough for the purposes of the simulation model. Also, we do not believe that goodness-of-fit tests (e.g., the chi-square test) are by themselves a definitive way of deciding how well a particular theoretical distribution fits a set of observed data. These tests have low power when the sample size n is "small," and may have high power (i.e., able to detect even small discrepancies between the fitted and underlying distributions) for "large" n . Thus, in the latter case, a theoretical distribution that is good enough for all practical purposes may be rejected because a large sample size happens to be available. We recommend using both tests and graphical comparisons to determine the efficacy of a particular theoretical distribution. For example, it is quite useful to plot a particular fitted density function (appropriately normalized) over a histogram and visually determine the quality of the "fit."

In some simulation studies it may not be possible to collect data on the random variables of interest, so the above techniques are not applicable to the problem of selecting corresponding probability distributions. For example, if the system being studied does not currently exist in some form, collecting data from the system is obviously not possible. This difficulty can also occur for existing systems, if the number of required probability distributions is large and the time available for the simulation study prohibits the necessary data collection and analysis. In many such situations, a good practical approach is to model an input random variable by a triangular distribution, which is parameterized by its minimum, maximum, and most likely values. People typically feel comfortable specifying these values and, furthermore, the triangular distribution is supported by almost all simulation software.

REFERENCE

Law, A.M. and W.D. Kelton (1991), *Simulation Modeling and Analysis*, Second Edition, McGraw-Hill, New York, NY.

BRUCE W. SCHMEISER

Broadly, families of input models can be evaluated via three criteria: (a) generality, (b) ease of generating realizations, and (c) ease of parameter selection. The importance of each differs according to the context in which the practitioner is working, which we discuss later. First we briefly discuss these three criteria.

Generality is the first criterion. Many classical input models are special or limiting cases of others. For example, the uniform distribution is a special case of the beta. The gamma distribution is a limiting case of the beta, and in turn the normal is a limiting case of the gamma and the exponential is a special case of the gamma. Thus the beta distribution never provides a worse input model than the gamma, normal, exponential, and uniform distributions. Generality can roughly be measured by the number of shape parameters; here the beta distribution has two, the gamma one, the normal none, the exponential none, and the uniform none. Schmeiser [1977] is an early review of general input models. Debrota et al. [1989] discuss ideas and software for fitting the four-parameter Johnson family of

distributions. The cost of the generality of the beta, and of other general families of distributions, lies with the other two criteria.

Random-variate generation, the second criterion, is related to input modeling in that some input models allow trivial, efficient random variate generation [Schmeiser 1980; Devroye 1986]. Ideally, inverse-transformation algorithms are also available to aid correlation induction for variance reduction. Schmeiser and Song [1989] look at inverse-distribution algorithms for a variety of well-known stochastic processes. Wilson and Avramidis [1989] look at general inverse-distribution-function input modes.

Parameter selection, the third criterion, arises in three ways. First is the classical situation of fitting the input model to real-world data; classical methods of fitting apply. The second is fitting the model to expert opinion, where ideally the input model has intuitive properties so that the opinion can be easily converted to parameter values. The third is selecting parameter values to provide a range of cases, no one of which is meant to match a particular real-world situation; rather the interest is in interpolating or optimizing among cases.

For discussion, we refer to the first two as *real-world modeling* and the third as *experimental-design modeling*. In the former, the true input processes are unknown and our input models are only reasonable approximations to the real world (e.g., the time between machine failures in an existing system). In the latter, the true input processes are assumed known, since the real world is one of assumption.

In experimental-design modeling the problem is to assume models that are tractable for determining parameter values so that we can easily select a set of cases (an experimental design) so that practitioners with real problems can interpolate between the cases (e.g., the power of a test of hypothesis for i.i.d. exponential data for sample size $n = 13$ when the cases run include $n = 10$ and $n = 15$). Fitting the input model to data is seldom an issue in experimental-design modeling. Monte Carlo experiments performed by statisticians fall in this category. By referring to these experiments as "experimental-design" modeling I do not mean to imply any particular formal or informal analysis of the results; the reference is only to the problem of choosing input models.

In contrast, real-world models are the domain of those modeling existing or proposed systems. Often no reason exists to expect that a good input model will be related to the classical distributions found in statistics textbooks. In addition, in real-world modeling the practitioner often has to fit many input models, usually in a short amount of time and often with little statistical background; the issue is then one of efficiently finding an adequate model rather than fine tuning to the model that would have been used given more time and training.

Another taxonomy of input models is by their dependency structure: (a) independent scalar distributions, (b) multivariate distributions, (c) scalar time-series models, and (d) point processes. This taxonomy is not perfect. For example, it does not include multivariate time series, a combination of (b) and (c), and point processes are often best viewed as the time-series of times between event epochs.

Most commercial software focuses on the classical scalar distributions. The first GPSS that I used supported the user only if the input model was normal, exponential, or uniform. The saving feature was that the user could specify a general piece-wise linear distribution function. But the ease of choosing a built-in distribution was tempting, even when all of the three built-in distributions were suspect. The current software is much better, but yet seldom provides built-in distributions beyond the classical scalar distributions.

A variety of models exist for multivariate models based on classical distributions. Schmeiser and Lal [1980] provide a brief survey; Johnson [1987] provides a comprehensive work on multivariate distributions for simulation in the experimental-design context, with emphasis on graphically depicting the distributions. Schmeiser and Lal [1982] discuss a family of bivariate gamma random vectors (and corresponding first-order autoregressive gamma time series) that will obtain any marginal gamma distributions and any corresponding feasible correlation, but the cost of such generality even in the restricted gamma case is solving nonlinear equations numerically. Lewis et al. [1989] discuss a variety of gamma time series.

In the real-world context random-vector input models can sometimes be constructed by cascading conditional distributions, often intuitively by building a logical model. In other cases, the conditional distributions are not intuitive. Then I like to transform multi-

variate normal random vectors to match desired marginal distributions and correlations, ideally with scatter plots to check that the correlation is capturing the desired dependence between components (see, e.g., Schmeiser [1990]). Similarly, ARMA time-series with normal marginal distributions can be transformed to nonnormal time series with similar dependency structures.

My final point concerns determining the validity of input models, an important issue in real-world modeling. Neyman-Pearson hypothesis testing is often advocated and sometimes used to assess model validity, with the null hypothesis being that the model is correct and the alternative hypothesis being that it is incorrect. But in real-world modeling the null hypothesis is almost certainly false, making model validation a strained application of the Neyman-Pearson approach. Such a test really is checking whether enough real-world data have been collected to detect the error that is present, not whether the model is valid. Since the null hypothesis is false, the real problem is to estimate the error in the input model and to determine whether it is negligible. Another way of saying this is that validation is not a matter of statistical significance, which is the focus of hypothesis testing, but of practical significance: whether the modeling error is acceptable in the situation at hand. Practical significance can often be assessed easily with simple graphs. For example, scalar marginal distributions can be assessed by plotting the real-world empirical distribution function against the fitted distribution function. (Another common practice is to plot two empirical distribution functions, the real-world data and a sample of equal size from the fitted input model. Plotting against the distribution function of the fitted model contains more information and therefore dominates.) Scatter plots and plots of time series are useful for input models with dependence.

REFERENCES

- DeBroya, D., R.S. Dittus, S.D. Roberts, J.R. Wilson, J.J. Swain, and S. Venkatraman (1989), "Modeling Input Processes with Johnson Distributions," In *Proceedings of the 1989 Winter Simulation Conference*, E.A. MacNair, K.J. Musselman, and P. Heidelberger, Eds., IEEE, Piscataway, NJ, 308-318.
- Devroye, L. (1986), *Non-Uniform Random Variate Generation*, Springer-Verlag, New York, NY.
- Johnson, M.E. (1986), *Multivariate Statistical Simulation*, Wiley, New York, NY.
- Schmeiser, B.W. (1977), "Methods for Modeling and Generating Probabilistic Components in Digital Computer Simulation when the Standard Distributions are not Adequate: A Survey," In *Proceedings of the 1977 Winter Simulation Conference*, 51-57. (Reprinted in *Simuletter* 10, Fall 1978/Winter 1978-79, 38-43, 72).
- Schmeiser, B.W. (1980), "Random Variate Generation: A Survey," In *Simulation with Discrete Models: A State of the Art View*, T.I. Oren, C.M. Shub, P.F. Roth, Eds., IEEE, Piscataway, NJ, 79-104.
- Schmeiser, B.W. (1990), "Simulation Experimentation," In *Handbook of Operations Research and Management Science, Volume 2: Stochastic Models*, D. Heyman and M. Sobel, Eds., North Holland, 295-330.
- Schmeiser, B.W. and R. Lal (1980), "Multivariate Modeling in Simulation: A Survey," In *33rd Annual Technical Conference Transactions*, American Society for Quality Control, 252-261.
- Schmeiser, B.W. and R. Lal (1982), "Bivariate Gamma Random Vectors," *Operations Research* 30, 355-374.
- Schmeiser, B.W. and W.-M.T. Song (1989), "Inverse-Transformation Algorithms for some Common Stochastic Processes," In *Proceedings of the 1989 Winter Simulation Conference*, E.A. MacNair, K.J. Musselman, and P. Heidelberger, Eds., IEEE, Piscataway, NJ, 490-496.
- Avramidis, T. and J.R. Wilson (1989), "A Flexible Method for Estimating Inverse Distribution Functions in Simulation Experiments," In *Proceedings of the 1989 Winter Simulation Conference*, E.A. MacNair, K.J. Musselman, and P. Heidelberger, Eds., IEEE, Piscataway, NJ, 428-436.

JAMES R. WILSON

In modeling and simulation of stochastic systems, a major problem is the selection of probability distributions that will adequately represent the input processes (populations) driving the simulation model. When it is feasible to collect sample data from a target population, simulation input modeling is usually accomplished by (a) hypothesizing a standard parametric distribution to describe that population, (b) estimating the associated parameters based on the sample information, and (c) performing diagnostic checks to assess the adequacy of the fit based on a comparison of the sample distribution with the fitted distribution. In the absence of sample information for parameter estimation and goodness-of-fit testing, practitioners usually try to elicit expert opinions about enough numerical characteristics of the target population (for example, the mode, the end points, or the mean) to specify uniquely a member of the hypothesized distribution family.

Rather than trying to fit a miscellaneous collection of distributions in an *ad hoc* fashion, it is preferable to adopt a systematic approach to input modeling based on a sufficiently general family of related distributions. This is hardly a new idea—it was the basic premise underlying the development of comprehensive distribution families by Pearson [1895] and Johnson [1949a, b]. In recent years the dramatic increase in the number of large-scale simulation studies has led to renewed interest in this approach [Ramberg and Schmeiser 1974; Schmeiser and Deutsch 1977; Johnson 1987; DeBroya et al. 1989a, b]. This discussion focuses on some of the main requirements of a general distribution family for simulation input modeling.

1. Flexibility. A great diversity of distributional shapes arise in modern simulation studies, and the overriding consideration in effective input modeling must be the flexibility of the distribution family to be used. The chief advantage of the Pearson and Johnson families is that they are capable of matching the first four moments of any distribution; and in many circumstances this provides sufficient flexibility to represent the important features of an empirical or theoretical distribution. The generalized lambda family of distributions [Ramberg and Schmeiser 1974] and the absolute lambda family [Schmeiser and Deutsch 1977] provide somewhat less flexibility, but all of these families encompass a much greater span of shapes than the standard distributions that are commonly used in simulation experiments.

2. Generalizability in one dimension. All of the distribution families discussed so far are based on univariate densities with at most four parameters. This limitation on the parameterization implies a corresponding limitation on the variety of distributional shapes that can be achieved; and it is not uncommon to encounter data sets in practice that simply cannot be adequately fitted by any of the well-known families of univariate distributions. The phrase *generalizability in one dimension* is used here to refer to the capability for modifying the basic functional form of a fitted univariate distribution by an open-ended, systematic generalization of the parameterization of that distribution. Hora [1983] and Avramidis and Wilson [1989] have proposed alternative techniques for achieving such generalizability, but it seems that many fundamental properties of these techniques have not been completely worked out; moreover there does not yet appear to be any substantial body of practical experience with either of these techniques.

3. Extendability to higher dimensions. This phrase refers to the ability to model not only univariate populations but also bivariate and higher-dimensional populations. Although the dependency structure among input variates is frequently ignored in practice, this can be a source of significant error in large-scale simulation studies [McDaniel et al. 1988]. In addition to its well-known flexibility, one of the principal advantages of the Johnson distribution family is that it can be naturally extended to fit multidimensional populations [Johnson 1949b; Johnson 1987].

4. Tractability. The ease with which a family can be manipulated analytically or numerically is an important consideration. The mathematically complex, highly ramified Pearson family has discouraged some potential users, and the Johnson family is also relatively difficult to handle. The absolute and generalized lambda families are substantially easier to handle, but they still present nontrivial mathe-

mathematical and numerical difficulties. The user can be shielded from this complexity by software packages that facilitate fitting distributions from a general family using sample data [Swain et al. 1988] or by subjective estimation [DeBroda et al. 1989]; nevertheless, the robustness and responsiveness of these packages can depend critically on the tractability of the underlying distribution family [AbouRizk et al. 1990].

5. Good parameterization. In addition to the mathematical and numerical tractability of a distribution family, another related but distinctly different consideration is the extent to which (a) the parameters of the constituent densities can be given direct physical interpretations, and (b) different parameters regulate genuinely different characteristics of the target population. For example, the parameters of Johnson distributions frequently are given a biometric interpretation in applications to forestry [Schreuder and Hafley 1977]. However, it is clear that sample estimators of the parameters of a Johnson distribution are highly correlated. Moreover, fits of nearly the same quality can frequently be obtained with a single Johnson density and with substantially different values of the parameters of that density. These properties of the Johnson family indicate that there is substantial overlap in the roles played by some of its parameters, especially in the bounded (S_B) subfamily of Johnson distributions. This is a disturbing state of affairs.

6. Ease of variate generation. Although this topic might be viewed as another aspect of tractability, it is considered separately for several reasons. In the first place, the variate-generation scheme should be compatible with the use of variance-reduction techniques [Wilson 1984]; and this suggests that the inverse transform method should always be used for sampling from univariate distributions. When sampling a multivariate distribution, it is highly desirable that each component of the generated random vector be a monotonically nondecreasing function of each random number required to generate the complete random vector. On the other hand, the efficiency with which random variates are generated does not appear to be a crucial issue in many large-scale simulation studies [Klein and Baris 1990]; thus the increased computational overhead required to use an appropriately monotonic random-variate generator should be more than compensated by the effective implementation of variance-reduction techniques like common random numbers and antithetic variates.

REFERENCES

- AbouRizk, S.M., D.W. Halpin, and J.R. Wilson (1990), "Visual Interactive Fitting of Beta Distributions," Technical Report SMS 90-13, School of Industrial Engineering, Purdue University, West Lafayette, IN.
- Avramidis, A.N. and J.R. Wilson (1989), "A Flexible Method for Estimating Inverse Distribution Functions in Simulation Experiments," In *Proceedings of the 1989 Winter Simulation Conference*, E.A. MacNair, K.J. Musselman, and P. Heidelberger, Eds., IEEE, Piscataway, NJ, 428-436.
- DeBroda, D.J., R.S. Dittus, S.D. Roberts, and J.R. Wilson (1989a), "Visual Interactive Fitting of Bounded Johnson Distributions," *Simulation* 52, 199-205.
- DeBroda, D.J., R.S. Dittus, J.J. Swain, S.D. Roberts, S. Venkatraman, and J.R. Wilson (1989b), "Modeling Input Processes with Johnson Distributions," In *Proceedings of the 1989 Winter Simulation Conference*, E.A. MacNair, K.J. Musselman, and P. Heidelberger, Eds., IEEE, Piscataway, NJ, 308-318.
- Hora, S.C. (1983), "Estimation of the Inverse Function for Random Variate Generation," *Communications of the Association for Computing Machinery* 26, 590-594.
- Johnson, M.E. (1987), *Multivariate Statistical Simulation*, Wiley, New York, NY.
- Johnson, N.L. (1949a), "Systems of Frequency Curves Generated by Methods of Translation," *Biometrika* 36, 149-176.
- Johnson, N.L. (1949b), "Bivariate Distributions Based on Simple Translation Systems," *Biometrika* 36, 297-304.
- Klein, R.W. and P.M. Baris (1990), "Selecting and Generating Variates for Modeling Service Times," *Computers and Industrial Engineering*, to appear.
- McDaniel, R.R., R.S. Sullivan, and J.R. Wilson (1988), "A Simulation Model for Welfare Policy Analysis," *Socioeconomic Planning Sciences* 22, 157-165.
- Pearson, K. (1895), "Contributions to the Mathematical Theory of Evolution.—II. Skew Variations in Homogeneous Material," *Philosophical Transactions of the Royal Society of London, Series A*, 186, 343-414.
- Ramberg, J.S. and B.W. Schmeiser (1974), "An Approximate Method for Generating Asymmetric Random Variables," *Communications of the Association for Computing Machinery* 17, 78-82.
- Schmeiser, B.W. and S.J. Deutsch (1977), "A Versatile Four Parameter Family of Probability Distributions Suitable for Simulation," *AIIE Transactions* 9, 176-182.
- Schreuder, H.T. and W.L. Hafley (1977), "A Useful Distribution for Describing Stand Structure of Tree Heights and Diameters," *Biometrics* 33, 471-478.
- Swain, J.J., S. Venkatraman, and J.R. Wilson (1988), "Least-Squares Estimation of Distribution Functions in Johnson's Translation System," *Journal of Statistical Computation and Simulation* 29, 271-297.
- Wilson, J.R. (1984), "Variance Reduction Techniques for Digital Simulation," *American Journal of Mathematical and Management Sciences* 4, 277-312.