Key words: *Monte Carlo, sampling, simulation, variance reduction*

An Illustration of the Sample Space Definition of Simulation and Variance Reduction

Barry L. Nelson
Department of Industrial and Systems Engineering
The Ohio State University
Columbus, Ohio 43210

ABSTRACT

A simulation experiment performed to estimate the expected number of customers served per day in a bank is used to illustrate a new definition of simulation experiments and a new taxonomy of variance reduction techniques. Particular emphasis is placed on how this framework might lead to automated variance reduction for general simulation experiments.

1. INTRODUCTION

Variance reduction techniques (VRTs) had their origins in survey sampling and Monte Carlo estimation, and there exist numerous VRTs and reasonably good guidelines for using them in these applications. Since random sampling always implies variability of estimators, it is natural, and frequently essential, to seek means of reducing variance. However, application of VRTs in simulation experiments is difficult because of the complicated, dynamic nature of simulation models. Research efforts have concentrated on the development of VRTs for narrow classes of simulation models, but there exist no guidelines for using VRTs in general simulation problems. Practitioners tend to view variance reduction as a complicated exercise with no guarantee of success. Automated variance reduction in standard simulation analysis packages is the only hope for widespread use of VRTs in practice, but a requirement for an automated procedure is a structured "world" in which to work.

Recently, Nelson and Schmeiser (1984a,b, 1985) developed a mathematical-statistical definition of simulation experiments and a taxonomy of VRTs based on this definition. The taxonomy views VRTs as compositions of transformations from six elemental classes, where a transformation maps one simulation experiment into

TRANSACTIONS of The Society for Computer Simulation
ISSN 0740-6797/85$2.00 + .10

Volume 2, Number 3, pp. 237–247

another experiment whose statistics may have smaller variance. This framework is intended to be a first step toward automated variance reduction.

In this paper, we illustrate certain aspects of the framework. The emphasis is on understanding the components of the "sample space definition," the six classes of transformations, and how this approach provides a new perspective on variance reduction that may eventually lead to automated procedures.

To make the illustration concrete, a textbook simulation experiment, found to be an excellent example, is used. The example (Bratley *et al.* 1983) describes a simulation experiment performed to estimate the expected number of customers served per day in a bank. The operation of the bank is modeled as follows:

> A bank employs three tellers. Each works at exactly the same rate and can handle exactly the same customers. Most days, all three tellers report for work. However, 15% of the time only two tellers are present, and 5% of the time only one teller turns up at work.
>
> The bank opens at 10 A.M., and closes at 3 P.M. From 9:45 A.M. until 11 A.M., and from 2 P.M until 3 P.M., one customer arrives every 2 minutes on average, and from 11 A.M. until 2 P.M. one customer arrives every minute. More precisely, the intervals between successive arrivals are independent and exponentially distributed with the appropriate mean. Those customers who arrive before 10 A.M. wait outside the door until the bank opens; at 3 P.M. the door is closed, but any customers already in the bank will be served.
>
> Customers form a single queue for the tellers. If, when a customer arrives, there are n people ahead of him in the queue (not counting the people receiving service), then he turns around and walks out ("balks" in queueing terminology) with the following probability:
>
> $$P[\text{balk}] = \begin{cases} 0 & n \leqslant 5 \\ (n-5)/5 & 6 \leqslant n \leqslant 9 \\ 1 & n \geqslant 10 \end{cases}$$
>
> The customer at the head of the queue goes to the first teller who is free. Customer service times are distributed according to an Erlang distribution with parameter 2 and mean service time 2 minutes. (Bratley *et al.*, 1983, pp. 13–14).

Although the model is simple, it is still too complicated for an analytic derivation of the expected number served. Bratley *et al.* (1983) illustrate simulation methodology (including a traditional treatment of variance reduction) using an experiment consisting of 200 simulated days of bank operation. Their text includes a FORTRAN program to execute the experiment. For the numerical results given later, we used their program and random number generator, and added a subroutine to initialize the simulation clock and event list. Only one random number stream was used, but the starting seeds for each random sequence (interarrival times, service times, number of tellers at work, and balking decision) were set 200,000 values apart. [See Bratley *et al.* (1983, p. 202) for the random number generator, p. 217 for the event calendar routines, and pp. 340–343 for the simulation code.]

2. A TAXONOMY OF VARIANCE REDUCTION TECHNIQUES

In the sample space definition of simulation experiments (Nelson and Schmeiser 1984a), a simulation experiment has a context (Ω, θ), where Ω is the sample space containing all possible realizations of the simulation input (see below) and θ is a vector of unknown parameters on interest. The sample space Ω represents all of the uncertain elements in the system, and θ is the performance measure that the simulation experiment is designed to estimate. The context defines the boundaries of the simulation experiment.

The simulation input, X, is a multivariate random variable modeled by the known cumulative probability distribution F over Ω; that is $X \sim F(x)$. Thus, we model the uncertain elements in the system by the probability distribution of their occurrence, and we will drive the experiment by generating realizations of X according to $F(x)$.

The simulation output, Y, is a function of the input and the sampling plan $R*$; that is $Y = g(X;R*)$. The function g embodies the logic of the system, in other words how it reacts to the uncertain elements. Thus, Y is the observed system behavior. The sampling plan does not correspond to any physical aspect of the system, but rather specifies a stopping rule for the execution of the experiment, since we could conceptually generate outputs indefinitely.

The statistic, Z, is a function of Y and is the estimator of θ; that is $Z = h(Y)$ and $E[Z] \cong \theta$. It is the variance of Z that VRTs reduce. Mathematically precise definitions of the components of the sample space definition can be found in Nelson and Schmeiser (1984a).

The bank simulation can be expressed in terms of the sample space definition as follows: The input, X, is composed of four sequences of random variables, the number of tellers on duty, the time between customer arrivals, an indicator of customer balking decisions, and customer service times. The probability distributions of these random variables are known and used to generate realizations, making them inputs. The output, Y, includes the actual arrival time of each customer, the number in line when each customer arrives, the time each customer completes service, and the number of customers served. The outputs are derived from the inputs, and are essential in the sense that all system performance measures can be derived from them (essential is a mathematical property needed to prove theoretical results; see Nelson and Schmeiser 1984a). The function g is implicit in the simulation code, but could conceptually be written explicitly. The sampling plan, $R*$, specifies that 200 observations of the daily number of customers served will be sampled. More generally, the sampling plan can specify the length of each output sequence. Finally, the statistic is initially the sample average of these 200 observations of daily number served.

Suppose Z and θ are scalars, and Z is an unbiased estimator of θ. Then the variance of Z can be written as

$$V[Z] = \int_{\Omega}[h(g(x;R*)) - \theta]^2 dF(x) \tag{1}$$

To reduce $V(Z)$, the components of the simulation must be redefined, or transformed, in a way that makes Eq. (1) smaller while preserving Z as an estimator of θ. In the sample space definition (Ω,θ) is considered fixed, so the reduction must be achieved by transforming F, g, $R*$ and/or h. Nelson and Schmeiser (1984b) define six mutually exclusive classes of transformations that exhaust the possible transformations under composition. In words, the classes can be loosely defined as follows:

1. Distribution Replacement (DR): Redefine the marginal distributions of the inputs without altering any statistical dependencies among the inputs.
2. Dependence Induction (DI): Redefine the statistical dependencies among the inputs without altering any marginal distributions of the inputs.
3. Equivalent Allocation (EA): Redefine the function g from input to output without altering the allocation of sampling effort.
4. Sample Allocation (SA): Redefine the allocation of sampling effort R_* without altering the function that defines the outputs.
5. Equivalent Information (EI): Redefine the function h from output to statistic without altering the argument of the statistic.
6. Auxiliary Information (AI): Redefine the argument of the statistic without altering the function from output to statistic.

VRTs are formed by composing members of these six classes of transformations. This is a radically different formulation of the variance reduction problem, but an appropriate one. Variance reduction is only achieved by redefining random variables in the simulation experiment. The sample space definition captures the definition of these random variables as they are described by the experimenter: either as uncertain elements sampled from known distributions, as the result of the reaction of the system to realizations of the uncertain elements, or as an aggregation of system performance that estimates the unknown system parameters. The six classes of transformations exhaust the ways these random variables can be redefined. For a very simple illustration of each class see Nelson and Schmeiser (1983), and for the decomposition of several well-known VRTs into their elemental transformations see Nelson and Schmeiser (1985).

There is one other key aspect of variance reduction. While there are many VRTs, not all of them will work in any particular experiment. In fact, a VRT applied in an inappropriate situation may actually increase variance, or worse may lead to an invalid estimator. How does one know if a VRT will be effective? The answer is prior knowledge, which we define as any information, either known with certainty or suspected, beyond what is needed to construct the original experiment. There are several sources of prior knowledge, including the experimenter, the model itself, or previous experimentation. The available prior knowledge is a real constraint on the variance reduction that can be achieved.

3. AUTOMATION

An automated variance reduction procedure based on our taxonomy might work in the following way (the framework presented here is part of ongoing research in conjunction with Bruce Schmeiser of Purdue University):

1. Express the simulation experiment in the standard form of the sample space definition.
2. Determine the available prior knowledge, based on a taxonomy of prior knowledge, from all available sources.

3. Match the prior knowledge to elemental transformations that require it to be effective.
4. Match the transformations to VRTs decomposed in terms of their elemental transformations.
5. Implement the VRT, perform the experiment, and evaluate the results.

Steps 2, 3 and 4 are likely to be iterative and interactive. We hasten to point out that more basic research is required before any such procedure can be implemented. A great deal of work has been done on step 4 (Nelson and Schmeiser 1985). However, the taxonomy of prior knowledge in step 2 is the subject of current research efforts. For the illustration below, we will use the following simple taxonomy that designates the source of prior knowledge as coming from

1. The experimenter's knowledge (PI.E)
2. The model itself (PI.M)
3. Universally true mathematical relationships (PI.U)
4. Pilot experimentation (PI.P)

The VRTs presented below are not new; we use familiar ones because our main purpose is to illustrate the taxonomy of variance reduction rather than survey VRTs. The taxonomy can also facilitate the discovery of new VRTs, as explained in Nelson and Schmeiser (1985). Because our perspective is radically different, we need to familiarize researchers and practitioners with it so that we can learn from their insights.

4. VARIANCE REDUCTION

The following notation will be used:

$Y_i \equiv$ number of customers served on day i
$\theta = E[Y_i] \equiv$ expected number of customers served on day i
$C(\ ,\ ) \equiv$ covariance of the enclosed quantities

In the original bank simulation the statistic is the sample average

$$Z = 200^{-1} \sum_{i=1}^{200} Y_i$$

The direct, or "crude," approach is to simulate 200 days of bank operation and use Z as the estimator of θ. Notice that the outputs from each day, $\{Y_i\}$, are independent and identically distributed random variables. The results of this simulation experiment, which match those of Bratley *et al.* (1983), are $Z = 240.92$ customers, with the estimated variance of Z being 2.78.

One statement of the variance reduction problem is to derive an unbiased estimator of θ that has significantly smaller variance than Z without expending too much additional effort to derive it. This is rather specific, and a more general characterization of variance reduction is given in Nelson and Schmeiser (1984b). All numerical results will be summarized in Section 5.

4.1. AI and EI

Combinations and transformations from EI and AI are the easiest to implement because they only redefine the statistic and do not alter the logic of the simulation model. For example, the experimenter may realize that the number of customers served each day differs from the number that arrive by the number that balk, an example of PI.E. Let A_i be the number of customers that arrive, and B_i be the number of customers that balk on day i, both auxiliary outputs that can be derived from the four essential outputs mentioned earlier. Then

$$\theta = E[Y_i] = E[A_i - B_i] = 247.5 - E[B_i]$$

which implies that θ can be indirectly estimated by estimating $E[B_i]$. Since A_i and B_i are both random variables, a source of variation is eliminated. The new statistic is

$$Z' = 247.5 - 200^{-1} \sum_{i=1}^{200} B_i \tag{2}$$

where the prior knowledge that $E[A_i] = 247.5$ comes from the known distribution of customer arrival times (PI.M). Equation (2) is an example of an indirect estimator (INDIR). INDIR tends to be a very problem specific strategy, depending almost entirely on the experimenter's knowledge. However, an automated procedure should permit the experimenter to incorporate such insights when they are available.

A more general strategy for using auxiliary information is to modify, or control, the outputs of interest based on knowledge of what the realization of an auxiliary output should have been. A key to such strategies is that the auxiliary output and the output of interest must be dependent random variables.

For example, it seems likely that the number of customers served each day in the bank depends on the number of tellers at work, since a longer queue of waiting customers increases the probability of balking for newly arriving customers. During the 200 days of simulated bank operation, we expect to see 160 days with three tellers at work, 30 days with two tellers, and 10 days when only one teller is present; this is known from the distribution of the number of tellers at work (PI.M). However, it is unlikely that the expected distribution will occur, which means that while our sample is random it is not entirely representative. In particular, the possible under- or overrepresentation of days with only one teller at work is a source of variability. This suggests trying to correct or control the sample by weighting each observation on the basis of whether the number of tellers was under- or overrepresented. This VRT is often called poststratified sampling, but we call it poststratifying the sample (PSTRAT) to make it clear that the sampling plan $R*$ is not affected.

An output is said to belong to stratum j if there were j tellers at work on the day the output was generated. Let Y_{ij} be the i^{th} such output, N_j the number of days with j tellers at work, $j = 1, 2, 3$, and $p_1 = .05$, $p_2 = .15$, and $p_3 = .80$. Then provided that there is at least one observation from each stratum

$$Z' = \sum_{j=1}^{3} \sum_{i=1}^{N_j} \frac{p_j}{N_j} Y_{ij} \tag{3}$$

is an unbiased estimator of θ (PI.U). The statistic Z gives each output weight 1/200, while Z′ gives weight p_j/N_j. If the outputs distribute themselves proportionately with respect to the strata (*i.e.*, $N_j = 200p_j$) then this reduces to 1/200. Otherwise, if stratum j is overrepresented the weight is less than 1/200, and the weight is greater than 1/200 if the stratum is underrepresented. The statistic Z′ uses the auxiliary information N_j to correct for disproportionate sampling; it also requires less specific prior knowledge than INDIR, and is thus more readily automatable.

Control variates (CV) are another means of correcting a nonrepresentative sample. The strategy is to use the difference between a parameter estimated from auxiliary outputs and its known expectation to adjust the estimator of the quantity of interest (θ, the daily number served in the example). Unlike PSTRAT, which only requires that the auxiliary output and the output of interest be dependent, CV requires a specific kind of dependence, namely linear correlation (Nelson 1985).

The expected number of customers to arrive each day is 247.5, as mentioned above. Clearly the number that arrives is closely related to the number that is served. In fact, if the queue remains short they should be almost identical. A common form of the CV estimator is the linear control, which in this example would be

$$Z' = Z - b(\bar{A} - 247.5) \tag{4}$$

where $\bar{A}$ is the sample mean of the $\{A_i\}$, and b is a constant that can be chosen to enhance the effectiveness of the control estimator.

The choice of b that minimizes the variance of Eq. (4) is $b^* = C(A_i, Y_i)/V(A_i)$ (see for instance Law and Kelton 1982). However, estimating b from the same simulation output can cause Eq. (4) to be biased. If the output of interest and the control variate are commensurate, meaning that they measure quantites in the same units, then setting $b = 1$ is sometimes an effective, though not optimal, strategy. It is clearly a strategy well-suited for automation. In our example Y_i and A_i both count numbers of customers, so $b = 1$ is an option.

A second approach is to estimate b^* from preliminary or pilot simulation experiments. This involves extra simulation effort, negating some of the gain from variance reduction. However, if we do not know a priori what auxiliary outputs might make good control variates (*i.e.*, which ones are strongly correlated with the outputs of interest), then pilot runs can supply this prior information (PI.P) as well as providing an independent estimate of b^*. An automated procedure should have the facility to perform pilot experiments when the necessary prior information is not available from other sources. It is also possible to induce dependence where it does not inherently exist using transformations in DR. Thus, one class of transformations can set up VRTs based on other classes.

Note that there are other techniques, such as Jackknifing (Bratley *et al.*, 1983), that deal with the bias from estimating b^* without the need of pilot runs. Depending on the situation, an automated procedure should be capable of choosing between the available options.

4.2. SA

When it is possible to fix the sampling plan SA strategies can be extremely effective. VRTs, such as stratified sampling (STRAT), allocate sampling effort deterministically to outputs that have specific characteristics. Unfortunately, in the simulation of complex, dynamic systems it is sometimes impossible to fix more than the overall simulation run duration. However, one notable exception is when the simulation involves independent replications with randomly selected initial conditions that can be fixed rather than randomly sampled. In our example, the number of tellers at work each day fits this criteria.

The question now becomes one of deciding the appropriate allocation of sampling effort. Let n_j be the number of days allocated with j tellers at work, $j = 1, 2, 3$. Note that n_j is part of the sampling plan $R*$, whereas N_j in PSTRAT was an auxiliary output. Proportional allocation, $n_j = 200p_j$, can be shown to be no worse than random sampling, and so is a good strategy for automation. However, given the strata the optimal allocation is (Cochran 1977)

$$n_j = 200 \frac{\sigma_j p_j}{\Sigma_{j=1}^{3} \sigma_j p_j}$$

where σ_j^2 is the variance of Y_i given there are j tellers at work. Notice that the optimal allocation is proportional not only to the probability of an observation from strata j (p_j), but also to the variance within the strata. The primary motivation behind VRTs based on SA is to allocate limited sampling effort where it does the most good. Since the σ_j are usually not known, they must be estimated to use optimal allocation. Here again, pilot runs can provide useful prior information.

The STRAT statistic is the same as the PSTRAT Eq. (3), but with N_j replaced by n_j. STRAT is generally more effective than PSTRAT, but requires control of the simulation sampling plan. PSTRAT is often referred to as a special case of STRAT, but the taxonomy reveals how they are actually different.

4.3. Combined Strategies

A taxonomy of VRTs that expresses VRTs as compositions of elemental transformations immediately suggests trying combinations of different classes. The difficulty for an automated procedure, and even for a sophisticated experimenter, is to recognize when the desired effect from applying one class of transformations conflicts with the effect from applying another. AI and EI are rather natural partners, since the first usually facilitates the second. SA transformations can sometimes be combined with these two.

In the bank example, we combined CV with PSTRAT and with STRAT by forming a CV estimator like Eq. (4) for each stratum individually, and then combining them. The new statistic is

$$Z' = \sum_{j=1}^{3} p_j\{\bar{Y}^j - b_j(\bar{A}^j - 247.5)\}$$

where $\bar{Y}^j$ and $\bar{A}^j$ are the sample means of $\{Y_i\}$ and $\{A_i\}$, respectively, from stratum j, based on N_j or n_j observations for PSTRAT or STRAT, respectively. STRAT and PSTRAT divide the simulation into three separate problems (days with 1, 2, or 3 tellers at work), to which CV can be applied individually. Thus, we do not expect a conflict. Another, unexplored possibility, is combining PSTRAT or STRAT with INDIR; again no conflict should result.

4.4. DR, DI and EA

Bratley *et al.* (1983) illustrate antithetic variates, which is based on DI transformations, and importance sampling, which is based on DR transformations, for the bank simulation; neither was effective. DI, while frequently used, is problematic because dependence induced between inputs may not yield the desired dependence between the outputs, unless we have a great deal of prior knowledge about how g maps inputs into outputs. VRTs based on DI are most useful for comparing alternative systems using common random numbers and to facilitate other VRTs as mentioned earlier. An illustration of DI for comparing alternatives is given in the next section. DR strategies are not well-suited for automation because, as Eq. (1) shows, the overall effect of distorting the input distribution is difficult to predict in a complex simulation.

EA transformations, not illustrated here, are sometimes used to set up strategies based on AI by producing new, auxiliary outputs. However, there are few VRTs based on EA because redefining g involves fundamentally redefining the logic of the system of interest. See Nelson and Schmeiser (1985) for an example.

5. NUMERICAL RESULTS SUMMARY

In this section we briefly summarize the results of simulation experiments performed using the VRTs illustrated in Section 4. These results are only intended to give a sense of the magnitude of variance reduction possible applying only simple VRTs that are suited to easy automation, and to demonstrate that intuition is not always a reliable guide for predicting the effectiveness of VRTs. Of course we can only estimate the variance reduction achieved, so our results are also subject to sampling variability.

The experiments all involved simulating 200 days of bank operation and employed the same random number streams as the crude experiment. Using the same streams is an example of DI, because the inputs of the various simulations are dependent, rather than being independently sampled. The sample space definition encompasses all nine simulations as one experiment, and the parameters of interest are the percentage variance reductions (which are differences). Positive dependence reduces the variance of estimators of relative differences and gives us more confidence that the estimated reductions below are valid.

In the cases of optimal CV and STRAT, 20 pilot runs (banking days) were made to estimate b and σ_j, respectively. We do not directly measure the additional effort required here since the purpose of this paper is illustration. However, we note that an automated procedure would have to aid the experimenter in making a decision about the effort to expend searching for and implementing a VRT.

The estimated variance for the PSTRAT experiment comes from a formula suggested by Cochran (1977), but modified for sampling from infinite populations. The formula, which is an approximation, is

$$V[Z'] = \sum_{j=1}^{3} \frac{p_j V[Y_{ij}]}{200} + \frac{(1-p_j)V[Y_{ij}]}{200^2}$$

Table 1 shows reductions ranging from a low of 39% to a high of 97%. More sophisticated strategies could almost certainly squeeze out further reductions. Of more interest here is that the combination of PSTRAT and CV, which employs only transformations in AI and EI, was extremely effective. With more and more simulation packages providing data bases for storing simulation outputs, such VRTs are particularly appealing.

Table 1
Experimental Results.

Experiment	Z	V(Z)	% Reduction
crude	240.92	2.78	—
INDIR	240.89	1.70	39%
PSTRAT	239.79	.92	67%
CV ($b^* = .83$)	240.54	1.73	38%
CV ($b = 1$)	240.46	1.70	39%
STRAT (optimal)	239.46	.92	69%
STRAT (prop.)	239.95	.99	64%
PSTRAT + CV (optimal)	239.11	.07	97%
STRAT + CV (optimal)	239.11	.08	97%

6. CONCLUSIONS

If VRTs are ever to be widely used in practice, variance reduction will have to be incorporated into standard simulation analysis packages as an automated, probably interactive procedure. To achieve automation, a structured environment such as the taxonomy of VRTs illustrated here is needed. While only a simple example has been used, more complicated simulation experiments introduce no additional complexity into the taxonomy.

Applying VRTs always involves a trade-off between the additional effort required and the variance reduction achieved. VRTs based on AI and EI, such as CV and PSTRAT, are good candidates for being totally automated, meaning applied without any help from the experimenter, because they only affect the statistical analysis and not the design of the simulation experiment. However, variance reductions of two or more orders of magnitude will usually require the experimenter's participation. Ways to include the experimenter, who may know little or nothing about variance reduction, in the procedure is still an open research question.

VRTs are often presented, as in this paper, in the context of simulating a single system. In this context the potential gains do not seem worth the effort. However, simulation experiments are more frequently performed to compare multiple system configurations, and the experimentation done sequentially. For example, when the simulation represents the objective function in an optimization problem, or when the results of the simulation are needed in real time. In these situations there is the potential for tremendous savings from variance reduction, and investing some effort in pilot experiments makes sense. There has yet to be any systematic study of how to effectively use pilot runs; *e.g.*, what runs should be made and what information should be extracted. Two examples were suggested above.

The key to applying VRTs is the availability of prior knowledge. The idea behind most VRTs is some mathematical relationship; *e.g.*, $V(X \pm Y) = V(X) + V(Y) \pm C(X,Y)$. Other sources of prior knowledge are the experimenter, the model, and pilot experiments. The illustration above showed contributions from all three. However, a taxonomy of prior knowledge that facilitates both discovering it and using it is needed to make automated variance reduction a reality.

REFERENCES

Bratley, P., B. L. Fox and L. E. Schrage. 1983. *A Guide to Simulation*. New York: Springer-Verlag.

Cochran, W. G. 1977. *Sampling Techniques*. New York: Wiley.

Law, A. M. and W. D. Kelton. 1982. *Simulation Modeling and Analysis*. New York: McGraw-Hill Book Co.

Nelson, B. L. and B. W. Schmeiser. 1983. "Variance Reduction: Basic Transformations," In *Proceedings of the 1983 Winter Simulation Conference* (J. Banks, S. Roberts, B. Schmeiser, eds.), IEEE, 255-258.

Nelson, B. L. and B. W. Schmeiser. 1984a. *A Mathematical-Statistical Framework for Variance Reduction, Part I: Simulation Experiments*, Research Memorandum No. 84-4, School of Industrial Engineering, Purdue University.

Nelson, B. L. and B. W. Schmeiser. 1984b. *A Mathematical-Statistical Framework for Variance Reduction, Part II: Classes of Transformations*, Research Memorandum No. 84-5, School of Industrial Engineering, Purdue University.

Nelson, B. L. and B. W. Schmeiser. 1985. "Decomposition of Some Well-Known Variance Reduction Techniques," *Journal of Statistical Computation and Simulation*. In press.

Nelson, B. L. 1985. *On Control Variate Estimators*, Working Paper Series No. 1985-006, Department of Industrial and Systems Engineering, The Ohio State University.

Received October 1984; accepted October 1985.

Barry L. Nelson is an Assistant Professor in the Department of Industrial and Systems Engineering at the Ohio State University. He is also associated with the Ohio State University Statistical Consulting Service. He obtained his Ph.D. from the School of Industrial Engineering at Purdue University. His research interests center on the design and analysis of computer simulation experiments on models of stochastic systems, particularly methods for statistically efficient simulation.

His continuing memberships include the ASA, IIE, ORSA, TIMS and SCS. Dr. Nelson is an active member of the TIMS College on Simulation and Gaming, and is assistant editor of its newsletter.