

## WHITE NOISE ASSUMPTIONS REVISITED: REGRESSION METAMODELS AND EXPERIMENTAL DESIGNS IN PRACTICE

Jack P. C. Kleijnen

Center for Economic Research and  
Department of Information Systems and Management  
Tilburg University  
5000 LE Tilburg, THE NETHERLANDS

### ABSTRACT

Classic linear regression metamodels and their concomitant experimental designs assume a univariate (not multivariate) simulation response and white noise. By definition, white noise is normally (Gaussian), independently (implying no common random numbers), and identically (constant variance) distributed with zero mean (valid metamodel). This advanced tutorial tries to answer the following questions: (i) How realistic are these classic assumptions in simulation practice? (ii) How can these assumptions be tested? (iii) If assumptions are violated, can the simulation's I/O data be transformed such that the analysis becomes correct? (iv) If such transformations cannot be applied, which alternative statistical methods (for example, generalized least squares, bootstrapping, jackknifing) can then be applied?

### 1 INTRODUCTION

Simulation models may be either deterministic or random (stochastic). To investigate the Input/Output (I/O) behavior of these simulation models, the analysts often use *linear regression* metamodels; for example, first-order and second-order polynomial approximations of the I/O function implied by the underlying simulation model. A good analysis (for example, a regression analysis) requires a good *statistical design*; for example, a fractional factorial such as a  $2^{k-p}$  design. For more mathematical details and background information I refer to my old textbook Kleijnen (1987) and my forthcoming textbook Kleijnen (2007); a recent tutorial is Kleijnen (2006).

In this article, I revisit the *classic assumptions* for linear regression analysis and its concomitant designs. These classic assumptions stipulate *univariate* output and *white noise*. In practice, however, these assumptions usually do not hold.

Indeed, in practice the simulation output (say)  $\hat{\Theta}$  is usually a *multivariate* random variable. For example, the simulation output (response)  $\hat{\Theta}_1$  may estimate the mean

flow time, and  $\hat{\Theta}_2$  may estimate the 90% quantile of the waiting time distribution. More examples will follow in Section 2.

*White noise* (say)  $u$  is Normally, Independently, and Identically Distributed (NIID) with zero mean:  $u \sim NIID(0, \sigma_u^2)$ . This definition implies the following assumptions:

1. *normally* (Gaussian) distributed simulation responses;
2. *no Common Random Numbers* (CRN) across the (say)  $n$  factor (input) combinations simulated;
3. a *common variance* (or homoscedasticity) of the simulation responses across these  $n$  combinations;
4. a *valid* regression metamodel; i.e., *zero* expected values for the residuals of the fitted metamodel.

In this article, I raise the following *questions*:

1. How *realistic* are these classic assumptions?
2. How can these assumptions be *tested* if it is not obvious that the assumption is violated (for example, if CRN are used, then the independence assumption is obviously violated)?
3. If an assumption is violated, can the simulation's I/O data be *transformed* such that the assumption holds?
4. If not, which *alternative statistical methods* can then be applied?

The answers to these questions are scattered throughout the literature on statistics and simulation. In this advanced tutorial, I therefore try to answer these questions in a coherent way. For more details (including additional references and examples) I refer to Kleijnen (2007).

The remainder of this article is organized as follows. Section 2 discusses multivariate simulation output. Section 3 addresses possible nonnormality of the simulation output, including tests of normality, transformations of simulation

I/O data, jackknifing, and bootstrapping. Section 4 covers variance heterogeneity (or heteroscedasticity) of the simulation output. Section 5 discusses cross-correlated simulation outputs, created through CRN. Section 6 discusses nonvalid low-order polynomial metamodells. Section 7 summarizes the major conclusions. A list with more than forty references enables the interested readers to study the topic of this tutorial, in much more detail.

## 2 MULTIVARIATE SIMULATION OUTPUT

### 2.1 Realistic Univariate Assumption?

In practice, the simulation model usually gives multivariate output. A class of practical examples concerns *inventory* simulation models with two outputs:

- (i) the sum of the holding and the ordering costs, averaged over the simulated periods;
- (ii) the service (or fill) rate, averaged over the same simulation periods.

The precise definitions of these costs and the service rate vary with the applications; see Law and Kelton (2000) and also Angün et al. (2006) and Ivanescu et al. (2006).

The *case study* in Kleijnen (1993) concerns a production planning Decision Support System (DSS), based on a simulation model. Originally, this simulation model had a multitude of outputs. However, to support decision making, it turned out that it sufficed to consider only the following two outputs (DSS criteria, bivariate response):

- (i) the total volume of steel tubes manufactured (which was of major interest to the production manager);
- (ii) the 90% quantile of delivery times (which was the sales manager's concern).

### 2.2 Multivariate Regression Metamodels

A *general notation* is

$$\mathbf{w} = s(d_1, \dots, d_k, \mathbf{r}_0) \quad (1)$$

with

$\mathbf{w}$ : vector of  $r$  simulation outputs, so  $\mathbf{w} = (w_0, \dots, w_{r-1})'$  (in simulation optimization it is traditional to label the  $r$  outputs starting with zero instead of one);

$s(\cdot)$ : mathematical function implicitly defined by the computer code implementing the given simulation model;

$d_j$ : factor (input variable)  $j$  of the simulation model, so  $\mathbf{D} = (d_{ij})$  is the design matrix for the simulation experiment, with  $j = 1, \dots, k$  and  $i = 1, \dots, n$  where  $n$  denotes the fixed number of combinations of the  $k$  factor levels (or values) in that experiment;

$\mathbf{r}_0$ : vector of PseudoRandom Number (PRN) seeds.

I assume that the multivariate I/O function  $s(\cdot)$  in Equation (1) is approximated by  $r$  univariate low-order

polynomials:

$$\mathbf{y}_h = \mathbf{X}\beta_h + \mathbf{e}_h \text{ with } h = 0, \dots, r-1 \quad (2)$$

with

$\mathbf{y}_h$ :  $n$ -dimensional vector  $(y_{1;h}, \dots, y_{n;h})'$  with the regression predictor  $y_h$  for simulation output  $w_h$ ;

$\mathbf{X}$ : common  $n \times q$  matrix of explanatory variables  $(\mathbf{x}_{ij})$  with  $\mathbf{x}_{ij}$  the value of explanatory variable  $j$  in combination  $i$  ( $i = 1, \dots, n; j = 1, \dots, q$ ); for simplicity, I assume that all fitted regression metamodels are polynomials of the same order (for example, either first order or second order) (if  $q > 2$  including an intercept, then the metamodel is called a *multiple* regression model);

$\beta_h$ :  $q$ -dimensional vector  $(\beta_{1;h}, \dots, \beta_{q;h})'$  with the  $q$  regression parameters for the  $h^{\text{th}}$  metamodel;

$\mathbf{e}_h$ :  $n$ -dimensional vector  $(e_{1;h}, \dots, e_{n;h})'$  with the residuals for the  $h^{\text{th}}$  metamodel, in the  $n$  combinations.

The  $r$  regression models in Equation (2) together have multivariate residuals  $\mathbf{e} = (\mathbf{e}_h)$  with the following two properties:

1. The univariate residuals  $\mathbf{e}_h = (e_{i;h})$  have variances that vary with the output variable  $w_h$  ( $h = 1, \dots, r$ ):  $\sigma_h^2 \neq \sigma^2$  (for example, simulated inventory costs and service percentages have different variances,  $\sigma_1^2 \neq \sigma_2^2$ ).
2. The univariate residuals  $e_{i;h}$  and  $e_{i;h'}$  are not independent:  $\sigma_{h;h';i} \neq 0$  for  $h \neq h'$ . Obviously, if these covariances (like the variances) would not vary with the combination  $i$ , then this property could be written as  $\sigma_{h;h';i} = \sigma_{h;h'} \neq 0$  for  $h \neq h'$  (for example, 'unusual' PRN streams in a given combination  $i$  may result in inventory costs that are 'relatively high'—that is, higher than expected—and a relatively high service percentage, so these two outputs are positively correlated:  $\sigma_{1;2} > 0$ ).

These two properties violate the classic assumptions. Consequently, it seems that the univariate Ordinary Least Squares (OLS) estimators should be replaced by the *Generalized Least Squares* (GLS) estimator of the parameter vector in the corresponding *multivariate regression* model. Fortunately, Rao (1959)—a more recent reference is Ruud (2000, p. 703)—proves that GLS reduces to OLS computed per output if the *same* matrix of independent variables is used (as is the case in Equation (2)); i.e., the Best Linear Unbiased Estimator is Estimator (BLUE) of  $\beta_h$  in Equation (2) is

$$\hat{\beta}_h = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{w}_h \text{ (} h = 0, \dots, r-1 \text{)} \quad (3)$$

where  $\mathbf{w}_h$  was defined below Equation (1), and  $\mathbf{D} = (d_{ij})$  defined below Equation (1) determines  $\mathbf{X}$  in Equations

(2) and (3). Given this result, the simulation analysts can easily obtain confidence intervals and statistical tests for the regression parameters per type of output variable; i.e., the analysts may indeed continue to use the classic formulas.

### 2.3 Designs for Multivariate Output

To the best of my knowledge, there are no *general* designs for multivariate output. Let's consider the following artificial example (inspired by Breukers 2006). The analysts are interested in two simulation response variables (so  $r = 2$  in Equation (1)). The number of simulation inputs is 15 (so  $k = 15$  in Equation (1)). First the analysts try to estimate the first-order effects, so they use a resolution-III design; for example, a  $2^{15-11}$  design (see Kleijnen 1987 and Kleijnen 2006). After running their simulation model with this design, they find that the factors labeled 1 through 7 have important main effects for response type 0, while the factors labeled 6 through 15 have important main effects for response type 1. In the next stage of their investigation, the analysts want to estimate the two-factor interactions between those factors that turned out to have important main effects in the first stage; i.e., the analysts use the *strong heredity* assumption in Wu and Hamada (2000), which states that if a factor has no important main effect, then this factor does not interact with any other factor. Because the number of two-factor interactions is  $k(k-1)/2$ , this number sharply increases with  $k$  (number of factors). In this example it is therefore efficient to estimate the interactions in two separate experiments, namely one experiment for each simulation response type. So the analysts split the original group of  $k = 15$  factors into one subgroup with  $k_0 = 7$  factors for simulation response 0, and  $k_1 = 10$  factors for the simulation response 1 (the factors 6 and 7 are members of both subgroups). The original group with 15 factors would require  $1+15+15 \times (15-1)/2 = 121$  factor combinations at least (a classic resolution-V design may require many more combinations, because these designs are often not saturated; see Kleijnen 1987 and Kleijnen 2006). Obviously, the first subgroup requires at least 29 combinations, and the second subgroup requires at least 56 combinations. So, together the two subgroups require at least  $29 + 56 = 85$  instead of 121 combinations; i.e., a 'divide and conquer' strategy pays off indeed.

## 3 NONNORMAL SIMULATION OUTPUT

Least Squares (LS) is a mathematical criterion, so LS does not assume a normal distribution. Only if the simulation analysts require statistical properties—such as BLUE, confidence intervals, and tests—then they usually assume a normal distribution. In this section, I try to answer the following questions (already formulated more generally in Section 1):

- (i) How realistic is the normality assumption?
- (ii) How can this assumption be tested?
- (iii) How can the simulation's I/O data be transformed such that the normality assumption holds?
- (iv) Which statistical methods can be applied that do not assume normality?

### 3.1 Realistic Normality Assumption?

By definition, *deterministic* simulation models do not have a normally distributed output for a given factor combination; this output is a single fixed value. In practice, simulation analysts often assume a normal distribution for the *residuals* of the fitted metamodel. An example is the case study in Kleijnen (1995) on coal mining using deterministic System Dynamics simulation; another example is the case study in Kleijnen, Van Ham, and Rotmans (1992) on global heating caused by the CO<sub>2</sub> greenhouse effect. Indeed, the simulation analysts might argue that so many things affect the residuals that the classic Central Limit Theorem (CLT) applies; i.e., a normal distribution is a good assumption for the residuals of a metamodel fitted to a deterministic simulation's I/O data.

In the remainder of this subsection, I focus on *random* simulation models. Simulation responses within a run are *autocorrelated* (serially correlated). By definition, a *stationary covariance process* has a constant mean (say  $E(w_t) = \mu$ ) and a constant variance  $var(w_t) = \sigma^2$ ; its covariances depend only on the *lag*  $|t - t'|$  between the variables  $w_t$  and  $w_{t'}$ ; that is,  $cov(w_t, w_{t'}) = \sigma_{|t-t'|}$ . The *average* of a stationary covariance process is *asymptotically* normally distributed if the covariances tend to zero sufficiently fast for large lags; see Lehmann (1999, Chapter 2.8). For example, in inventory simulations the output is often the costs averaged over the simulated periods; this average is probably normally distributed. Another output of an inventory simulation may be the service percentage calculated as the fraction of demand delivered from on-hand stock per (say) week, so 'the' output is the average per year computed from these 52 weekly averages. This yearly average may be normally distributed—unless the service goal is 'close' to 100%, so the average service rate is cut off at this threshold and the normal distribution is a bad approximation.

Note that confidence intervals based on Student's  $t$  statistic are quite insensitive to nonnormality, whereas the lack-of-fit  $F$ -statistic is more sensitive to nonnormality; see Kleijnen (1987) for details including references.

In summary, a limit theorem may explain why random simulation outputs are asymptotically normally distributed. Whether the actual simulation run is long enough, is always hard to know. Therefore it seems good practice to check whether the normality assumption holds (see the next subsection).

### 3.2 Testing Normality

Basic statistics textbooks (also see the recent article, Arcones and Wang 2006) and simulation textbooks (for example, Kleijnen 1987 and Law and Kelton 2000) propose several *visual plots* and *goodness-of-fit statistics* to test whether a set of observations come from a specific distribution type such as a normal distribution. A basic assumption is that these observations are IID. Simulation analysts may therefore obtain ‘many’ (say,  $m = 100$ ) replicates for a specific factor combination (for example, the base scenario) if such an approach is computationally feasible. However, if a single simulation run takes relatively much computer time, then only ‘a few’ (say,  $2 \leq m \leq 10$ ) replicates are feasible, so the plots are too rough and the goodness-of-fit tests lack power.

Actually, the white noise assumption concerns the meta-model’s *residuals*  $e$ —not the simulation model’s outputs  $w$ ; see Equation (2). The estimated residuals are  $\hat{e}_i = \hat{y}_i - w_i$  with  $i = 1, \dots, n$  and  $\hat{y}_i = \mathbf{x}_i \hat{\beta}$ ; an alternative definition is  $\hat{e}_i = \hat{y}_i - \bar{w}_i$  where  $\bar{w}_i = \sum_{r=1}^{m_i} w_{i,r} / m_i$  is the simulation output averaged over the  $m_i$  replicates. I assume that the simulation analysts obtain at least a few replicates,  $m_i > 1$ . For simplicity of presentation, I further assume that the number of replicates is constant:  $m_i = m (> 1)$ . If the simulation outputs  $w$  have a constant variance ( $\sigma_w^2$ ), then  $\sigma_w^2 / m$  ( $= \sigma_w^2 / m$ ) is also constant. Unfortunately, even if the average simulation outputs have a constant variance ( $\sigma_w^2 / m$ ) and are independent (no CRN), the *estimated* residuals do not have a constant variance and they are not independent; it can be proven that

$$\text{cov}(\hat{\mathbf{e}}) = [\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}']\sigma_w^2 \quad (4)$$

where  $\mathbf{X}$  is the  $n \times q$  matrix of explanatory regression variables defined below Equation (3). Nevertheless, analysts (for example, Ayanso, Diaby, and Nair 2006) apply visual inspection of residual plots, which are standard output of many statistical packages. Note that Equation (4) uses the well-known *hat matrix*  $\mathbf{H} = \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$ ; also see Atkinson and Riani (2000).

### 3.3 Transformations, Jackknifing, and Bootstrapping

The simulation output  $w$  may be transformed to obtain better normality. Well-known is the *Box-Cox power transformation*:

$$v = \frac{w^\lambda - 1}{\lambda} \text{ if } \lambda \neq 0; \text{ else } v = \ln(w). \quad (5)$$

A complication is that the metamodel now explains not the behavior of the original output, but the behavior of the transformed output! See Atkinson and Riani (2000, p. 82) and Freeman and Modarres (2006).

In case of nonnormal output, *outliers* occur more frequently when the actual distribution has ‘fatter’ tails. *Robust regression analysis* might then be applied; see Atkinson and Riani (2000) and Salibian-Barres (2006). However, I have not seen any applications of this approach in simulation.

Normality is not assumed by the following two general computer-intensive statistical procedures that use the original simulation I/O data ( $\mathbf{D}, \mathbf{w}$ ): jackknifing and bootstrapping (actually, the jackknife is a linear approximation of the bootstrap; see Efron and Tibshirani 1993). Both procedures have become popular since powerful and cheap computers have become available to the analysts.

#### 3.3.1 Jackknifing

In general, *jackknifing* solves the following two types of problems:

1. How to compute *confidence intervals* in case of nonnormal observations?
2. How to reduce possible *bias* of estimators?

Examples of nonnormal observations are the estimated service rate close to 100% in inventory simulations, and extreme quantiles such as the 99.99% point in risk simulations (see the nuclear waste simulations in Kleijnen and Helton 1999). Examples of biased estimators will follow in Section 4.

Suppose the analysts want a confidence interval for the regression coefficients  $\beta$  in case the simulation output has a very nonnormal distribution. So the linear regression metamodel is still Equation (2), but now with  $r = 1$ . Assume that each factor combination  $i$  is replicated an equal number of times,  $m_i = m > 1$ . The original OLS estimator (also see Equation (3)) is

$$\hat{\beta} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\bar{\mathbf{w}}. \quad (6)$$

Jackknifing deletes the  $r^{\text{th}}$  replicate among the  $m$  IID replicates, and recomputes the estimator for which a confidence interval is wanted:

$$\hat{\beta}_{-r} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\bar{\mathbf{w}}_{-r} \quad (r = 1, \dots, m) \quad (7)$$

where  $\bar{\mathbf{w}}_{-r}$  is the  $n$ -dimensional vector with elements that are the averages of the  $m - 1$  replicates after deleting replicate  $r$ :

$$\bar{w}_{i,-r} = \frac{\sum_{r' \neq r}^m w_{i,r'}}{m - 1} \quad (8)$$

where for the case  $r = m$  the summation runs from 1 to  $m - 1$  (not  $m$ ).

Obviously, Equation (7) gives the  $m$  correlated (vector) estimators  $\hat{\beta}_{-1}, \dots, \hat{\beta}_{-m}$ . For ease of presentation, I focus

on  $\beta_q$  (the last of the  $q$  individual regression parameters in the vector  $\beta$ ). Jackknifing uses the *pseudovalue* (say)  $J$ , which is the following weighted average of  $\widehat{\beta}_q$  (the original estimator) and  $\widehat{\beta}_{q;-r}$  (the  $q^{\text{th}}$  element of the jackknifed estimator  $\widehat{\beta}_{-r}$  defined in Equation (7)) with the number of observations as weights:

$$J_r = m\widehat{\beta}_q - (m-1)\widehat{\beta}_{q;-r}. \quad (9)$$

In this example both the original and the jackknifed estimators are unbiased, so the pseudovalues also remain unbiased estimators. Otherwise it can be proven that the bias is reduced by the *jackknife point estimator*

$$\bar{J} = \frac{\sum_{r=1}^m J_r}{m}, \quad (10)$$

which is simply the average of the  $m$  pseudovalues defined in Equation (9).

To compute a *confidence interval*, jackknifing treats the pseudovalues as if they were NIID:

$$P(\bar{J} - t_{m-1;1-\alpha/2}\widehat{\sigma}_{\bar{J}} < \beta_q < \bar{J} + t_{m-1;1-\alpha/2}\widehat{\sigma}_{\bar{J}}) = 1 - \alpha \quad (11)$$

where  $t_{m-1;1-\alpha/2}$  denotes the  $1 - \alpha/2$  quantile (upper  $\alpha/2$  point) of the distribution of Student's  $t$  statistic with  $m - 1$  degrees of freedom, and

$$\widehat{\sigma}_{\bar{J}} = \sqrt{\frac{\sum_{r=1}^m (J_r - \bar{J})^2}{m(m-1)}}.$$

The interval in Equation (11) may be used to test the null-hypothesis that the true regression parameter has a specific value, such as zero.

*Applications* of jackknifing in simulation are numerous. For example, jackknifing gave confidence intervals for Weighted LS (WLS) with estimated covariance matrix  $\widehat{\text{cov}}(\mathbf{w})$ ; see Kleijnen et al. (1987) and Section 4. Jackknifing reduced the bias and gave confidence intervals for a Variance Reduction Technique (VRT) called control variates or regression sampling; see Kleijnen et al. (1989). Jackknifing may also be applied in the renewal analysis of steady-state simulation (renewal analysis uses ratio estimators, which are biased); see Kleijnen and Van Groenendaal (1992, pp. 202-203).

### 3.3.2 Bootstrapping

Textbooks on bootstrapping are Davison and Hinkley (1997), Efron and Tibshirani (1993), Good (2005), and Lunneborg (2000); a recent article is Davidson and MacKinnon (2006)—more references will follow below. Bootstrapping may be used to solve two types of problems:

1. The relevant distribution is not Gaussian.
2. The statistic is not standard.

*Sub 1:* Reconsider the example used for jackknifing; i.e., the analysts want a confidence interval for the regression coefficients  $\beta$  in case of nonnormal simulation output. Again assume that each of the  $n$  factor combinations is replicated an equal number of times,  $m_i = m > 1$  ( $i = 1, \dots, n$ ). The original LS estimator was given in Equation (6).

The bootstrap distinguishes between the *original observations*  $w$  and the *bootstrapped observations* (say)  $w^*$  (note the superscript). Standard bootstrapping assumes that the original observations are IID. In the example, there are  $m_i = m$  IID original simulated observations per factor combination  $i$ , namely  $w_{i;1}, \dots, w_{i;m}$  (these observations give  $\bar{w}_i$ , which give the vector  $\bar{\mathbf{w}}$ , which occurs in Equation (6)).

The bootstrap observations are obtained by *resampling with replacement* from the original observations, while the sample size is kept constant, at  $m$ . In the example, the bootstrapped observations  $w_{i;1}^*, \dots, w_{i;m}^*$  occur with frequencies  $f_1, \dots, f_m$  such that  $f_1 + \dots + f_m = m$ ; i.e., these frequencies follow the multinomial (or polynomial) distribution with parameters  $m$  and  $p_1 = \dots = p_m = 1/m$ . This resampling is executed for each combination  $i$  ( $i = 1, \dots, n$ ). These bootstrapped outputs  $w_{i;1}^*, \dots, w_{i;m}^*$  give the bootstrapped average simulation output  $\bar{\mathbf{w}}^*$ . Substitution into Equation (6) gives the bootstrapped LS estimator

$$\hat{\beta}^* = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\bar{\mathbf{w}}^*. \quad (12)$$

To reduce sampling variation, this resampling is repeated (say)  $B$  times;  $B$  is known as the *bootstrap sample size* (typical values for  $B$  are 100 and 1,000). This gives  $\hat{\beta}_1^*, \dots, \hat{\beta}_B^*$  (or  $\hat{\beta}_b^*$  with  $b = 1, \dots, B$ ).

Let's again focus on the single regression parameter,  $\beta_q$ . The bootstrap literature gives several confidence intervals, but most popular is

$$P(\widehat{\beta}_{q;(\lfloor B\alpha/2 \rfloor)}^* < \beta_q < \widehat{\beta}_{q;(\lfloor B(1-\alpha/2) \rfloor)}^*) = 1 - \alpha$$

where  $\widehat{\beta}_{q;(\lfloor B\alpha/2 \rfloor)}^*$  is the  $\alpha/2$  quantile of the Empirical Density Function (EDF) of the bootstrap estimate  $\widehat{\beta}_q^*$ , and  $\widehat{\beta}_{q;(\lfloor B(1-\alpha/2) \rfloor)}^*$  is its  $1 - \alpha/2$  quantile.

*Applications* of bootstrapping include Kleijnen, Cheng, and Bettonvil (2001), validating trace-driven simulation models in case of serious nonnormal outputs (the test statistic was the difference between the average output of the real and the simulated systems).

*Sub 2:* Besides classic statistics such as the  $t$  and  $F$  statistics, the simulation analysts may be interested in statistics that have no tables with critical values, which provide confidence intervals—assuming normality. For example,

Kleijnen and Deflandre (2006) bootstrapped  $R^2$  to test the validity of regression metamodels in simulation.

#### 4 HETEROSCEDASTIC SIMULATION OUTPUT

By definition, *deterministic* simulation models give a single fixed value for a given factor combination, so the conditional variance is zero:  $\text{var}(w|\mathbf{x}) = 0$ . Simulation analysts often assume a normal distribution for the residuals of the metamodel fitted to the I/O data of the deterministic simulation model (see Section 3.1). Usually, the analysts then assume a normal distribution with a *constant* variance (Kriging models also assume a constant variance). I do not know a better assumption that works in practice for deterministic simulation models.

I further focus on *random* simulation models, and try to answer the following questions:

1. How *realistic* is the common variance assumption?
2. How can this assumption be *tested*?
3. How can the simulation's I/O data be *transformed* such that the common variance assumption holds?
4. Which statistical *analysis* methods can be applied that allow nonconstant variances?
5. Which statistical *design* methods can be applied to account for variance heterogeneity?

##### 4.1 Realistic Constant Variance Assumption?

In practice, the variances of random simulation outputs change when factor combinations change. For example, in the M/M/1 queueing simulation not only the mean of the steady-state waiting time changes as the traffic rate changes—the variance of this output changes even more! So the constant variance assumption seems very unrealistic!

##### 4.2 Testing the Constant Variance Assumption

Though it may be *a priori* certain that the variances of the simulation outputs are not constant, the analysts may hope that the variances are (nearly) constant in their particular application. Unfortunately, the variances are unknown so they must be estimated. This estimator itself has high variance; in case of normally distributed output,  $\text{var}(\hat{\sigma}^2) = 2\sigma^4/m$ . Actually, there are  $n$  combinations of the  $k$  factors in the simulation experiment, so  $n$  variance estimators  $\hat{\sigma}_i^2$  need to be compared. This problem may be solved in many different ways, but I recommend the distribution-free test in Conover (1980, p. 241).

##### 4.3 Variance Stabilizing Transformation

The logarithmic transformation in Equation (5) may be used not only to obtain normal output but also to obtain outputs

with constant variances. A problem may again be that the metamodel now explains the transformed output instead of the original output.

#### 4.4 Alternative Analysis Methods

In case of heterogeneous variances, the LS criterion still gives an *unbiased* estimator (it suffices that the residuals have zero mean,  $E(\mathbf{e}) = 0$ ). The variance of the LS (or OLS) estimator, however, now is

$$\text{cov}(\hat{\beta}) = (\mathbf{X}'_N \mathbf{X}_N)^{-1} \mathbf{X}'_N \text{cov}_N(\mathbf{w}) \mathbf{X}_N (\mathbf{X}'_N \mathbf{X}_N)^{-1} \quad (13)$$

where I explicitly denote the number of rows  $N = \sum_{i=1}^n$  of  $\mathbf{X}_N$ , which is an  $N \times q$  matrix;  $\text{cov}_N(\mathbf{w})$  is an  $N \times N$  matrix, and the first  $m_1$  elements on its main diagonal are  $\text{var}(w_1)$ , ..., the last  $m_n$  elements on this main diagonal are  $\text{var}(w_n)$ . In Section 5, I shall present a simple method to derive confidence intervals for the  $q$  individual OLS estimators  $\hat{\beta}_j$  (see Equation (24)).

Though the OLS estimator remains unbiased, it is no longer the BLUE. It can be proven that the BLUE is now the Weighted LS (or WLS) estimator

$$\tilde{\beta} = (\mathbf{X}'_N [\text{cov}(\mathbf{w})]^{-1} \mathbf{X}_N)^{-1} \mathbf{X}'_N [\text{cov}(\mathbf{w})]^{-1} \mathbf{w}. \quad (14)$$

For a constant number of replicates ( $m_i = m$ ), the WLS estimator may also be written as

$$\tilde{\beta} = (\mathbf{X}' [\text{cov}(\bar{\mathbf{w}})]^{-1} \mathbf{X})^{-1} \mathbf{X}' [\text{cov}(\bar{\mathbf{w}})]^{-1} \bar{\mathbf{w}} \quad (15)$$

where  $\mathbf{X}$  is  $n \times q$  and  $\text{cov}(\bar{\mathbf{w}}) = \text{cov}(\mathbf{w})/m$  where  $\text{cov}(\mathbf{w})$  is  $n \times n$ . The covariance matrix of this WLS estimator is

$$\text{cov}(\tilde{\beta}) = (\mathbf{X}' [\text{cov}(\bar{\mathbf{w}})]^{-1} \mathbf{X})^{-1}. \quad (16)$$

In practice,  $\text{cov}(\mathbf{w})$  is unknown so this covariance matrix must be estimated. The elements on this diagonal matrix are estimated through the classic unbiased variance estimator

$$\widehat{\text{var}}(w_i) = \hat{\sigma}^2(w_i) = s_i^2(w) = \frac{\sum_{r=1}^m (w_{ir} - \bar{w}_i)^2}{m-1}, \quad (17)$$

which gives  $\widehat{\text{cov}}(\mathbf{w})$ . Substituting this estimated matrix into the classic WLS formula Equation (14) gives the *Estimated WLS* (EWLS) or Aitken estimator. For a constant number of replicates this EWLS estimator is

$$\widehat{\tilde{\beta}} = (\mathbf{X}' [\widehat{\text{cov}}(\bar{\mathbf{w}})]^{-1} \mathbf{X})^{-1} \mathbf{X}' [\widehat{\text{cov}}(\bar{\mathbf{w}})]^{-1} \bar{\mathbf{w}}. \quad (18)$$

*This EWLS is not a linear estimator.* Consequently, the statistical analysis becomes more complicated. For example, the analogue of Equation (16) holds only asymptotically

(under certain conditions); see, for example, Godfrey (2006) and Kleijnen, Cremers, and Van Belle (1985):

$$\mathbf{cov}(\hat{\beta}) \approx (\mathbf{X}' [\mathbf{cov}(\bar{\mathbf{w}})]^{-1} \mathbf{X})^{-1}. \quad (19)$$

Classic confidence intervals no longer hold.

Relatively simple solutions for this type of problem were already presented in Sections 3.3.1 and 3.3.2, namely jackknifing and bootstrapping. *Jackknifing* the EWLS estimator was done in Kleijnen et al. (1987), as follows. Delete the  $r^{th}$  replicate among the  $m$  IID replicates, and recompute the EWLS estimator (analogous to Equation (7)):

$$\hat{\beta}_{-r} = (\mathbf{X}' [\widehat{\mathbf{cov}}(\bar{\mathbf{w}})_{-r}]^{-1} \mathbf{X})^{-1} \mathbf{X}' [\widehat{\mathbf{cov}}(\bar{\mathbf{w}})_{-r}]^{-1} \bar{\mathbf{w}}_{-r}$$

where  $\bar{\mathbf{w}}_{-r}$  consists of  $n$  averages computed from  $m - 1$  replicates after deleting replicate  $r$ , and  $\widehat{\mathbf{cov}}(\bar{\mathbf{w}})_{-r}$  is computed from the same replicates. These  $\hat{\beta}_{-r}$  and the original  $\hat{\beta}$  computed through Equation (18) can be used to compute the pseudovalues, which give the desired confidence interval. *Bootstrapping* the EWLS estimator is discussed in Kleijnen and Deflandre (2006).

#### 4.5 Alternative Designs

If the output variances are not constant, *classic designs* still give the *unbiased* OLS estimator  $\hat{\beta}$  and WLS estimator  $\hat{\beta}$ . The literature pays little attention to the derivation of alternative designs for heterogeneous output variances. Kleijnen and Van Groenendaal (1995) investigated designs in which the  $n$  factor combinations are replicated so many times that the estimated variances of the averages per combination are (approximately) constant. Because  $var(\bar{w}_i) = \sigma_i^2/m_i$  ( $i = 1, \dots, n$ ), the number of replicates should satisfy

$$m_i = c_0 \sigma_i^2 \quad (20)$$

where  $c_0$  is a common positive constant such that the  $m_i$  become integers. This equation implies that the higher the variability of the simulation output  $w_i$  is, the more replicates are simulated. The allocation of the total number of simulation runs ( $N = \sum_{i=1}^n m_i$ ) according to Equation (20) is not necessarily optimal, but it simplifies the regression analysis and the design of the simulation experiment (an alternative design replaces  $\sigma_i^2$  by  $\sigma_i$ ). Indeed the regression analysis can now apply OLS to the averages  $\bar{w}_i$  to get BLUE.

In practice, however, the variances of the simulation outputs must be estimated. A *two-stage* procedure takes a *pilot sample* of (say)  $m_0 \geq 2$  replicates for each factor

combination, and estimates the variances  $\sigma_i^2$  through

$$s_i^2(m_0) = \frac{\sum_{r=1}^{m_0} [w_{ir} - \bar{w}_i(m_0)]^2}{m_0 - 1} \quad (i = 1, \dots, n) \quad (21)$$

with  $\bar{w}_i(m_0) = \sum_{r=1}^{m_0} w_{ir}/m_0$ . Combining Equations (21) and (20), Kleijnen and Van Groenendaal (1995) selects additional replicates  $\hat{m}_i - m_0$  where

$$\hat{m}_i = m_0 \left\lceil \frac{s_i^2(m_0)}{\min_{1 \leq i \leq n} s_i^2(m_0)} \right\rceil$$

with  $\lceil x \rceil$  denoting the integer closest to  $x$  (so, in the second stage no additional replicates are simulated for the combination with the smallest estimated variance). After the second stage all  $\hat{m}_i$  replicates are used to estimate the average output and its variance. OLS is applied to these averages. The covariance matrix  $\mathbf{cov}(\hat{\beta})$  is estimated through Equation (13) with  $\mathbf{cov}(\bar{\mathbf{w}})$  estimated through a diagonal matrix with diagonal elements  $s_i^2(\hat{m}_i)/\hat{m}_i$ . Confidence intervals are based on the classic  $t$  statistic with degrees of freedom equal to  $m_0 - 1$ .

Because these  $s_i^2(\hat{m}_i)/\hat{m}_i$  may still differ considerably, this two-stage approach may be replaced by a *sequential* approach. The latter approach adds one replicate at a time, until the estimated variances of the average simulation outputs have become practically constant; see Kleijnen and Van Groenendaal (1995). The sequential procedure requires fewer simulation responses, but is harder to understand, program, and implement.

## 5 COMMON RANDOM NUMBERS

In this section, I try to answer the following questions:

1. How *realistic* is it to assume the use of CRN?
2. Which statistical *analysis* methods can be applied that allow CRN?
3. Which statistical *design* methods can be applied to account for CRN?

### 5.1 Realistic CRN Assumption?

Obviously, CRN are applied in *random* simulation only. In practice, simulation analysts often use CRN; actually, CRN is the *default* of much simulation software. The goal of CRN is to reduce  $var(\hat{\beta}_j)$  with  $j = 1, \dots, q$  (actually, the variance of the intercept increases when CRN are used). So CRN is useful to better explain the factor effects, and to better predict the output of combinations not yet simulated (provided the inaccuracy of the estimated intercept is outweighed by the accuracy of all other estimated effects). The use of CRN creates correlation between the simulation outputs  $w_{i,r}$  and  $w_{i',r}$  ( $i, i' = 1, \dots, n; r = 1, \dots, m; m = \min_i m_i$ ).

## 5.2 Alternative Analysis Methods

So, CRN violates the classic assumptions of regression analysis. The analysts have two options (analogous to the options in case of heterogeneous variances):

- (i) Continue to use OLS
- (ii) Switch to GLS.

*Sub (i):* The variance of the OLS estimator is given by Equation (13), but now  $\text{cov}(\mathbf{w})$  is not a diagonal matrix. This covariance matrix may be estimated by  $\widehat{\text{cov}}(\mathbf{w})$  with the elements

$$\widehat{\text{cov}}(w_i, w_{i'}) = \frac{\sum_{r=1}^m (w_{i;r} - \bar{w}_i)(w_{i';r} - \bar{w}_{i'})}{m-1} \quad (22)$$

where  $m = \min(m_i, m_{i'})$ ; usually,  $m_i = m_{i'} (= m)$  if CRN is applied. This  $\widehat{\text{cov}}(\mathbf{w})$  is *singular* if the number of replicates is ‘too small’; that is, if  $m \leq n$ ; see Dykstra (1970).

Kleijnen (1992) showed that confidence intervals for the  $q$  individual OLS estimators can be computed from a Student  $t$  statistic with  $m-1$  degrees of freedom—provided  $m > n$  (in this statistic, the standard errors  $s(\hat{\beta}_j)$  are the square roots of the elements on the main diagonal of the corrected covariance matrix in Equation (13)).

An *alternative* method does not require the estimation of  $\text{cov}(\bar{\mathbf{w}})$  to derive confidence intervals for the OLS estimators, so it suffices that  $m > 1$ . This alternative requires  $m$  computations of the OLS estimator (also see Law and Kelton 2000, p. 630, 642); i.e., from replicate  $r$  the analysts compute

$$\hat{\beta}_r = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{w}_r \quad (r = 1, \dots, m). \quad (23)$$

The  $n$  elements of the vector  $\mathbf{w}_r$  are correlated (because they use CRN) and may have different variances, but the  $m$  estimators  $\hat{\beta}_{j;r}$  of a specific regression parameter  $\beta_j$  are independent (because they use non-overlapping PRN streams) and have a common standard deviation (say)  $\sigma(\hat{\beta}_j)$ . So

$$t_{m-1} = \frac{\bar{\hat{\beta}}_j - \beta_j}{s(\bar{\hat{\beta}}_j)} \quad \text{with } j = 1, \dots, q \quad (24)$$

with

$$s(\bar{\hat{\beta}}_j) = \sqrt{\frac{\sum_{r=1}^m (\hat{\beta}_{j;r} - \bar{\hat{\beta}}_j)^2}{m(m-1)}}.$$

*Sub (ii):* CRN implies that the BLUE is the GLS estimator; see Equation (14) where  $\text{cov}(\mathbf{w})$  is now not diagonal. Obviously,  $\text{cov}(\hat{\beta})$  is analogous to Equation (16). Substituting  $\widehat{\text{cov}}(\mathbf{w})$  into the classic GLS formula gives the *Estimated GLS* (EGLS), which is the analogue of the

EWLS estimator in Equation (18). The EGLS estimator can again be analyzed through jackknifing and bootstrapping. Kleijnen (1992) compared OLS and EGLS relying on the asymptotic covariance matrix Equation (19) with nondiagonal  $\widehat{\text{cov}}(\mathbf{w})$ ; Davidson and MacKinnon (2006), however, claims that ‘bootstrap tests ... yield more reliable inferences than asymptotic tests in a great many cases’.

In conclusion, CRN with EGLS may give better point estimates of the factor effects than CRN with OLS, but the EGLS estimate requires ‘many’ replicates—namely  $m > n$ —to obtain a nonsingular  $\widehat{\text{cov}}(\mathbf{w})$ .

## 5.3 Alternative Designs for CRN

The literature pays no attention to the derivation of alternative designs for CRN. *Sequential* procedures are proposed in Kleijnen and Van Beers (2004) and Van Beers and Kleijnen (2006). These designs select the next factor combination to be simulated, where the simulation model may be either deterministic or random. The simulation I/O data ( $\mathbf{D}, \mathbf{w}$ ) are now analyzed through Kriging (instead of linear regression), which allows the simulation outputs of different combinations to be correlated (as Kriging assumes a stationary covariance process).

## 6 NONVALID LOW-ORDER POLYNOMIAL METAMODEL

Now, I try to answer the following questions:

1. How can the validity of the low-order polynomial metamodel be tested?
2. If this metamodel is not valid, how can the simulation’s I/O data be transformed such that a low-order polynomial becomes valid?
3. Which alternative metamodels can be applied?

### 6.1 Testing the Validity Variance Assumption

A valid metamodel has zero mean residuals, so  $H_0 : E(e) = 0$ . To test this null-hypothesis, the analysts may apply the classic *lack-of-fit F-statistic* assuming white noise. However, if the analysts apply CRN, then they may apply Rao’s variant derived in Rao (1959) (and evaluated in Kleijnen 1992):

$$F_{n-q; m-n+q} = \frac{m-n+q}{(n-q)(m-1)} (\bar{\mathbf{w}} - \hat{\mathbf{y}})' [\widehat{\text{cov}}(\bar{\mathbf{w}})]^{-1} (\bar{\mathbf{w}} - \hat{\mathbf{y}}) \quad (25)$$

where  $n > q$ ,  $m > n$ , and  $\hat{\mathbf{y}}$  denotes the EGLS estimator. Obviously, this test also allows EWLS instead of EGLS. Normality of the simulation output is an important assumption for both the classic test and Rao’s test. In case of nonnormality, the analysts may apply jackknifing or boot-



strapping; Kleijnen and Deflandre (2006) bootstraps Rao's statistic and the classic  $R^2$  statistic.

An alternative test uses *cross-validation* and the  $t$  statistic, which is less sensitive to nonnormality than the  $F$  statistic; see Kleijnen (1992). Moreover, this  $t$  statistic requires fewer replications, namely  $m > 1$  instead of  $m > n$  if EWLS or EGLS is used.

Besides these quantitative tests, the analysts may use *graphical* methods to judge the validity of a fitted metamodel (be it a linear regression model or some other type of metamodel such as a Kriging model). Scatterplots are well known; they are the standard output of many software packages. The recent panel publication Simpson et al. (2004) also emphasizes the importance of visualization; also see Helton et al. (2006).

If a validation test rejects  $H_0$ , then the analysts may consider the following alternatives.

## 6.2 Transformations

A well-known transformation in queueing simulations combines two simulation inputs—namely, the arrival rate  $\lambda$  and the service rate  $\mu$ —into a single independent regression variable—namely, the traffic rate  $x = \lambda/\mu$ . Another transformation replaces  $y$ ,  $\lambda$ , and  $\mu$  by  $\log(y)$ ,  $\log(\lambda)$ , and  $\log(\mu)$ , to make the first-order polynomial approximate relative changes.

Another simple transformation assumes that the I/O function of the underlying simulation model is *monotonic*. Then the dependent and independent variables may be replaced by their ranks, which results in so-called *rank regression*; see Conover and Iman (1981) and Saltelli and Sobol (1995). Note that Spearman's correlation coefficient uses the same transformation for two correlated random variables. For example, Kleijnen and Helton (1999) applies rank regression and Spearman's coefficient to find the most important factors in a simulation model of nuclear waste disposal.

Transformations may also be applied to make the simulation output (dependent regression variable) better satisfy the assumptions of normality (see Equation (5)) and variance homogeneity. Unfortunately, different goals of the transformation may conflict with each other; for example, the analysts may apply the logarithmic transformation to reduce nonnormality, but this transformation may give a metamodel in variables that are not of immediate interest.

If classic low-order polynomial metamodels (estimated from their concomitant designs) do not give valid metamodels, then I recommend to look for transformations, as discussed above. I do not recommend routinely adding higher-order terms to the metamodel, because these terms are hard to interpret. However, if the goal is not to better *understand* the underlying simulation model but to better *predict* the output of an expensive simulation model, then

high-order terms may be added. Indeed, full factorial  $2^k$  designs enable the estimation of all interactions (for example, the interaction among all  $k$  factors).

If more than two levels are simulated per factor, then the following types of metamodels may be considered.

## 6.3 Alternative Metamodels

There are several alternative metamodel types; for example, Kriging models. These alternatives may give better predictions than low-order polynomials do. However, these alternatives are so complicated that they do not help the analysts better understand the underlying simulation model—except for sorting the simulation inputs in order of their importance. Furthermore, these alternative metamodels require *alternative design types* (for example, Latin Hypercube Sampling and max-min designs). This is a completely different issue, so I refer to the extensive literature on this topic (including Kleijnen 2007).

## 7 CONCLUSIONS

In this advanced tutorial, I discussed the assumptions of classic linear regression analysis and the concomitant statistical designs when these methods are applied in simulation practice. In Section 2, I pointed out that multivariate simulation output can still be analyzed through OLS. In Section 3, I addressed possible nonnormality of simulation output, including normality tests, transformations of simulation I/O data, jackknifing, and bootstrapping. In Section 4, I presented analysis and design methods for heteroscedastic simulation output. In Section 5, I discussed how to analyze simulation I/O data that uses CRN, so the simulation outputs are correlated (across different factor combinations, but within the same replicate). In Section 6, I discussed possible lack-of-fit tests for low-order polynomial metamodels, transformations to improve the metamodel's validity, and alternative metamodels and designs. Throughout this paper, I gave many references for further study of these issues.

## REFERENCES

- Angün, E., D. den Hertog, G. Gürkan, and J.P.C. Kleijnen. 2006. Response surface methodology with stochastic constrains for expensive simulation. Working Paper, Tilburg University, Tilburg, Netherlands.
- Arcones, M. A., and Y. Wang. 2006. Some new tests for normality based on U-processes. *Statistics & Probability Letters* 76(1): 69-82.
- Atkinson, A., and M. Riani. 2000. *Robust diagnostic regression analysis*. New York: Springer.
- Ayanso, A., M. Diaby, and S. K. Nair. 2006. Inventory rationing via drop-shipping in Internet retailing: a

- sensitivity analysis. *European Journal of Operational Research* 171(1): 135-152.
- Breukers, A. 2006. Bio-economic modelling of brown rot in the Dutch potato production chain. Doctoral dissertation, Wageningen University, Wageningen, The Netherlands.
- Conover, W. J. 1980. *Practical nonparametric statistics*, second edition. New York: Wiley.
- Conover, W. J., and R. L. Iman. 1981. Rank transformations as a bridge between parametric and nonparametric statistics. *The American Statistician* 35(3): 124-133.
- Davidson, R., and J. G. MacKinnon. 2006. Improving the reliability of bootstrap tests with the fast double bootstrap. *Computational Statistics & Data Analysis*, in press.
- Davison, A. C., and D. V. Hinkley. 1997. *Bootstrap methods and their application*. Cambridge: Cambridge University Press.
- Dykstra, R. L. 1970. Establishing the positive definiteness of the sample covariance matrix. *The Annals of Mathematical Statistics* 41(6): 2153-2154.
- Efron, B., and R. J. Tibshirani. 1993. *An introduction to the bootstrap*. New York: Chapman & Hall.
- Freeman, J., and R. Modarres. 2006. Inverse Box Cox: the power-normal distribution. *Statistics & Probability Letters* 76(8): 764-772.
- Godfrey, L. G. 2006. Tests for regression models with heteroskedasticity of unknown form. *Computational Statistics & Data Analysis* 50(10): 2715-2733.
- Good, P. I. 2005. *Resampling methods: a practical guide to data analysis*, third edition. Boston: Birkhäuser.
- Helton, J. C., J. D. Johnson, C. J. Sallaberry, and C. B. Storlie. 2006. Survey of sampling-based methods for uncertainty and sensitivity analysis. *Reliability Engineering and Systems Safety*, in press.
- Ivanescu, C., W. Bertrand, J. Fransoo, and J. P. C. Kleijnen. 2006. Bootstrapping to solve the limited data problem in production control: an application in batch processing industries. *Journal of the Operational Research Society* 57(1): 2-9.
- Kleijnen, J. P. C. 1987. *Statistical tools for simulation practitioners*. New York: Marcel Dekker.
- Kleijnen, J. P. C. 1992. Regression metamodelling for simulation with common random numbers: comparison of validation tests and confidence intervals. *Management Science* 38(8): 1164-1185.
- Kleijnen, J. P. C. 1993. Simulation and optimization in production planning: a case study. *Decision Support Systems* 9: 269-280.
- Kleijnen, J. P. C. 1995. Sensitivity analysis and optimization of system dynamics models: regression analysis and statistical design of experiments. *System Dynamics Review* 11(4): 275-288.
- Kleijnen, J. P. C. 2006. Regression models and experimental designs: a tutorial for simulation analysts. CentER Discussion Paper, number 2006-10.
- Kleijnen, J. P. C. 2007. *DASE: design and analysis of simulation experiments*. Springer Science + Business Media.
- Kleijnen, J. P. C., R. C. H. Cheng, and B. Bettonvil. 2001. Validation of trace-driven simulation models: bootstrapped tests. *Management Science* 47(11): 1533-1538.
- Kleijnen, J. P. C., P. Cremers, and F. van Belle. 1985. The power of weighted and ordinary least squares with estimated unequal variances in experimental designs. *Communications in Statistics, Simulation and Computation* 14(1): 85-102.
- Kleijnen, J. P. C., and D. Deflandre. 2006. Validation of regression metamodelling in simulation: bootstrap approach. *European Journal of Operational Research* 170(1): 120-131.
- Kleijnen, J. P. C., and J. Helton. 1999. Statistical analyses of scatter plots to identify important factors in large-scale simulations, 1: review and comparison of techniques. *Reliability Engineering and Systems Safety* 65(2): 147-185.
- Kleijnen, J. P. C., P. C. A. Karremans, W. K. Oortwijn, and W. J. H. van Groenendaal. 1987. Jackknifing estimated weighted least squares: JEWLS. *Communications in Statistics, Theory and Methods* 16(3): 747-764.
- Kleijnen, J. P. C., J. Kriens, H. Timmermans, and H. Van den Wildenberg. 1989. Regression sampling in statistical auditing: a practical survey and evaluation (including Rejoinder). *Statistica Neerlandica* 43(4): 193-207 (225).
- Kleijnen, J. P. C., and W. C. M. van Beers. 2004. Application-driven sequential designs for simulation experiments: Kriging metamodelling. *Journal of the Operational Research Society* 55(9): 876-883.
- Kleijnen, J. P. C., and W. van Groenendaal. 1992. *Simulation: a statistical perspective*. Chichester (England): John Wiley.
- Kleijnen, J. P. C., and W. van Groenendaal. 1995. Two-stage versus sequential sample-size determination in regression analysis of simulation experiments. *American Journal of Mathematical and Management Sciences* 15(1&2): 83-114.
- Kleijnen, J. P. C., G. van Ham, and J. Rotmans. 1992. Techniques for sensitivity analysis of simulation models: a case study of the CO<sub>2</sub> greenhouse effect. *Simulation* 58(6): 410-417.
- Law, A. M., and W. D. Kelton. 2000. *Simulation modeling and analysis*, third edition. Boston: McGraw-Hill.
- Lehmann, E. L. 1999. *Elements of large-sample theory*. New York: Springer.

- Lunneborg, C. E. 2000. *Data analysis by resampling: concepts and applications*. Pacific Grove, California: Duxbury Press.
- Rao, C. R. 1959. Some problems involving linear hypothesis in multivariate analysis. *Biometrika* 46: 49-58.
- Rao, C. R. 1967. Least squares theory using an estimated dispersion matrix and its application to measurement of signals. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability* 1, 355-372.
- Ruud, P. A. 2000. *An introduction to classical econometric theory*. New York: Oxford University Press.
- Salibian-Barrera, M. 2006. Bootstrapping MM-estimators for linear regression with fixed designs. *Statistics & Probability Letters*, in press.
- Saltelli, A., and I.M. Sobol. 1995. About the use of rank transformation in sensitivity analysis of model output. *Reliability Engineering and System Safety* 50: 225-239.
- Simpson, T. W., A. J. Booker, D. Ghosh, A. A. Giunta, P. N. Koch, and R.-J. Yang. 2004. Approximation methods in multidisciplinary analysis and optimization: a panel discussion. *Structural and Multidisciplinary Optimization* 27(5): 302-313.
- Van Beers, W. C. M., and J. P. C. Kleijnen. 2006. Customized sequential designs for random simulation experiments: Kriging metamodeling and bootstrapping. Working Paper, Tilburg University, Tilburg, Netherlands.
- Wu, C. F. J., and M. Hamada. 2000. *Experiments; planning, analysis, and parameter design optimization*. New York: Wiley.

## AUTHOR BIOGRAPHY

**JACK P.C. KLEIJNEN** is professor of ‘simulation and information systems’ at Tilburg University. He is a member of the Department of Information Management and the Operations Research Group of the Center for Economic Research (CentER). He also teaches at the Technical University Eindhoven, in the Postgraduate International Program in Logistics Management Systems. He is an ‘external fellow’ of the Mansholt Graduate School of Social Sciences of Wageningen University. His research concerns the statistical design and analysis of simulation experiments, information systems, and supply chains. He has been a consultant for several organizations in the USA and Europe, and serves on many international editorial boards and scientific committees. He spent several years in the USA, at universities and private companies. He received a number of international awards, including the INFORMS Simulation Society’s ‘Lifetime Professional Achievement Award (LPAA)’ of 2005. More information can be found on <http://center.uvt.nl/staff/kleijnen/>.