

SENSITIVITY ANALYSIS AND OPTIMIZATION IN SIMULATION: DESIGN OF EXPERIMENTS AND CASE STUDIES

Jack P.C. Kleijnen

Department of Information Systems and Auditing/Center for Economic Research (CentER)
School of Management and Economics
Tilburg University (Katholieke Universiteit Brabant)
5000 LE Tilburg, NETHERLANDS

ABSTRACT

This paper is an advanced tutorial on the use of statistical techniques in sensitivity analysis, including the application of these techniques to optimization and validation of simulation models. Sensitivity analysis is divided into two phases. The first phase is a pilot stage, which consists of screening or searching for the important factors; a simple technique is sequential bifurcation. In the second phase, regression analysis is used to approximate the input/output behavior of the simulation model. This regression analysis gives better results when the simulation experiment is well designed, using classical statistical designs such as fractional factorials. To optimize the simulated system, Response Surface Methodology (RSM) is applied; RSM combines regression analysis, design of experiments, and steepest ascent. To validate a simulation model that lacks input/output data, again regression analysis and design of experiments are applied. Several case studies are summarized; they illustrate how in practice statistical techniques can make simulation studies give more general results, in less time.

1 INTRODUCTION

The *objective* of this paper is to examine the problem of sensitivity analysis in simulation, including the related issues of optimization and validation. To solve these problems, this contribution gives a *survey* of certain statistical techniques, namely Design Of Experiments (DOE) and its analysis through regression analysis (also known as ANOVA, ANalysis Of VAriance).

This paper is an advanced tutorial, which discusses not only methodology, but also applications. The reader is assumed to have a basic knowledge of mathematical statistics and simulation.

More specifically, the following *questions* are addressed (which should be answered for *all* simulation models):

1. *What if*: what happens if the analysts change parameters, input variables or modules (such as subroutines for priority rules) of the simulation model? This question is closely related to sensitivity analysis and optimization, as we shall see. The literature on DOE uses the term *factor* to denote a parameter, input variable or module.

2. *Validation*: is the simulation model an adequate representation of the corresponding system in the real world? This paper addresses only part of the validation problem.

To answer these practical questions, this paper takes techniques from the *science of mathematical statistics* (briefly, statistics). It is not surprising that statistics is so important in simulation: by definition, *simulation* means that a model is 'solved' by experimentation. But experimentation requires a good design and a good analysis! DOE with its concomitant analysis is a standard topic in statistics. However, the standard statistical techniques must be adapted such that they account for the *particularities of simulation*. For example, there are a great many factors in many practical simulation models. Indeed, one application (discussed later) has hundreds of factors, whereas standard DOE assumes only up to (say) fifteen factors. Moreover, stochastic simulation models use pseudorandom numbers, which means that the analysts have much more control over the noise in their experiments than the investigators have in standard statistical applications (for example, common and antithetic seeds may be used; see Kleijnen and Van Groenendaal 1992).

The *main conclusions* of this paper will be:

(i) *Screening* may use the simple, efficient, and effective technique of sequential bifurcation; see Bettonvil and Kleijnen (1994).

(ii) Next, *regression analysis* generalizes the results of the simulation experiment, since it characterizes the input/output behavior of the simulation model.

(iii) *Statistical designs* give good estimators of main (first-order) effects and interactions among factors; these designs require fewer simulation runs than intuitive designs do.

(iv) *Optimization* may use RSM, which builds on regression analysis and DOE; see (ii) and (iii).

(v) *Validation* may use regression and DOE, especially if there are no data on the input/output of the simulation model or its modules.

(vi) These statistical techniques have already been *applied* many times in practical simulation studies, in many domains; these techniques make simulation studies give more general results, in less time.

The remainder of this paper is organized as follows. §2 discusses sensitivity analysis by means of DOE, which treats the simulation model as a black box. More specifically, §2.1 studies the screening phase of a simulation study: which factors among the many potentially important factors are really important? A sub-subsection (§2.1.1) discusses a very efficient screening technique, called sequential bifurcation. §2.2 discusses how to approximate the input/output behavior of simulation models by regression analysis. First it discusses graphical methods, namely scatter plots; see §2.2.1. Next it presents regression analysis (which formalizes the graphical approach), including standardization of factors, Generalized Least Squares (GLS), and cross-validation; see §2.2.2. Next, §2.3 discusses statistical designs. First the focus is on designs that assume only main effects (§2.3.1). Then follow designs that give unbiased estimators for the main effects, even if there are interactions between factors (§2.3.2). Further, this section discusses designs that allow estimation of individual interactions (§2.3.3). §2.3 ends with designs for estimating the curvature (quadratic effects) of the input/output approximation (§2.3.4). §3 proceeds with the role of sensitivity analysis in validation, emphasizing the effects of data availability. §4 presents the optimization of simulated systems through RSM. §5 gives a summary and conclusions. Seventeen references conclude the paper. This paper is based on Kleijnen (1996).

2 SENSITIVITY ANALYSIS

The vast literature on simulation does not provide a standard definition of sensitivity analysis. This paper defines sensitivity analysis as the systematic investigation of the reaction of the simulation responses to *extreme* values of the model's input or to drastic chan-

ges in the model's structure. For example, what happens to a queueing simulation's response, when the customer arrival rate doubles; what happens if the priority rule changes from First In First Out (FIFO) to Shortest Processing Time (SPT)? So the focus is not on marginal changes (or perturbations) in the input values.

Moreover, the simulation model is treated as a *black box*: the simulation inputs and outputs are observed, and from this input/output behavior the factor effects are estimated. This approach is standard in DOE.

DOE has *advantages and disadvantages*. One benefit is that this approach can be applied to all simulation models. A drawback is that it can not take advantage of the specific structure of a given simulation model, so it may take many simulation runs to perform the sensitivity analysis. But DOE requires fewer runs than the naive approach often followed in practice (as we shall see).

Note: The intricacies of the specific simulation model at hand are considered in perturbation analysis and in modern importance sampling, also known as score function; see Ho and Cao (1991), Glynn and Iglehart (1989), and Rubinstein and Shapiro (1993) respectively. Perturbation analysis and score function require only one run. Unfortunately, these methods also require more mathematical sophistication.

2.1 Pilot or Screening Phase

In the pilot phase of a simulation study there are usually *a great many* potentially important factors. For example, in a queueing model with m servers there are m service rates; the queueing priority rule may also be a factor. It is the mission of science to come up with a short list of the most important factors; it is unacceptable to say 'everything depends on everything else': *parsimony* principle.

In *practice*, analysts often restrict their study to a few factors, usually no more than ten. Those factors are selected through intuition, prior knowledge, and the like. The factors that are ignored (kept constant), are - explicitly or implicitly- assumed to be unimportant. For example, in the simulation of queueing networks, it is traditional to assume equal service rates. Of course, such an assumption severely restricts the generality of the simulation study!

The *statistics literature* includes screening designs. These designs provide scientific methods for finding the important factors. There are several types of screening designs: random, supersaturated, group screening designs, and so on; see Kleijnen (1987).

Unfortunately, the statistics literature pays too little attention to screening designs. The reason for this neglect is that in standard statistical applications it is virtually impossible to control hundreds of factors; fifteen is hard enough. In *simulation*, however, models may have hundreds of parameters, and yet their control is simple: just specify which combinations of parameter values to simulate. Nevertheless, screening applications in simulation are rare, because most analysts are not familiar with these designs. Yet these designs are simple and efficient.

Recently, screening designs have been improved and new variations have been developed; details are given in Bettonvil and Kleijnen (1994). The next subsection discusses the most promising type, namely sequential bifurcation.

2.1.1 Sequential Bifurcation

Sequential bifurcation uses the *aggregation* principle, which is often applied in science when studying complicated systems. So at the start of the simulation experiment, sequential bifurcation groups the individual factors into clusters. To make sure that individual factor effects do not cancel out, sequential bifurcation assumes that the analysts know whether a specific individual factor has a positive or negative effect on the simulation response: *known signs*. In practice this assumption is not very restrictive. For example, in the queueing example it is known that increasing the service rates decreases waiting times (but it is unknown how big this decrease is; therefore the analysts use a simulation model).

In practice, sequential bifurcation was applied to an ecological simulation with 281 parameters. The ecological experts felt comfortable specifying in which direction a specific parameter affects the response (this response is the future carbon-dioxide or CO₂ concentration; CO₂ creates the greenhouse effect). Moreover, if a few individual factors have unknown signs, then these factors can be investigated separately, outside the sequential bifurcation design.

Sequentialization means that factor combinations to be simulated, are selected as the experimental results become available; that is, as simulation runs are executed, insight into factor effects is accumulated and used to select the next run. As the experiment proceeds, groups of factors are eliminated, because sequential bifurcation concludes that these clusters contain no important factors.

Also, as the experiment proceeds, the groups become smaller. More specifically, each group that seems to include one or more important factors, is split into

two subgroups of the same size: *bifurcation*. At the end of screening by means of sequential bifurcation, individual factors are investigated.

In the ecological application, sequential bifurcation took 154 simulation runs to identify and estimate the 15 most important factors among the original 281 factors. Moreover, had the analysts assumed no interactions between factors, then the number of runs could have been halved ($154/2 = 77$ runs).

The ecological case study concerns a deterministic simulation model (consisting of a set of non-linear difference equations). There is a need for more research, applying sequential bifurcation to large random simulations, such as simulations of queueing networks in telecommunications.

2.2 Approximating the Input/Output Behavior of Simulation Models by Regression Analysis

2.2.1 Introduction: Graphical Methods

After the screening phase (§2.1), the number of factors to be further investigated is reduced to a small number (for example, fifteen).

Practitioners often make a *scatter plot* with on the x -axis the values of one factor (for example, service rate) and on the y -axis the simulation response (say, average waiting time). This graph indicates the input/output behavior of the simulation model, treated as a black box. It shows whether this factor has a positive or negative effect on the response, whether that effect remains constant over the domain (experimental area) of the factor, etc.

This scatter plot can be further analyzed: *fit a curve* to these (x, y) data; for example, fit a straight line ($y = \beta_0 + \beta_1 x$). Of course, other curves can be fitted: quadratic (second degree polynomial), exponential, logarithmic (using paper with a log scale), and so on.

To study *interactions* between factors, scatter plots per factor can be combined. For example, the scatter plot for different service rates was drawn, given a certain number of servers. Plots for different numbers of servers can now be superimposed. Intuitively, the curve for a low number of servers lies above the curve for a high number of servers (if not, the simulation model is probably wrong; see the discussion on validation in §3). If the response curves are not parallel, there are interactions, by definition.

However, superimposing many plots is cumbersome. Moreover, their interpretation is subjective: are the response curves really parallel straight lines? These shortcomings are removed by regression analysis.

2.2.2 Regression analysis

A regression model is a *metamodel* of the simulation model; that is, a regression model approximates the input/output behavior of the simulation model that generates the input/output data to which the regression analysis is applied. Consider the *second degree polynomial*

$$Y_i = \beta_0 + \sum_{h=1}^k \beta_h x_{ih} + \sum_{h=1}^k \sum_{h'=1}^k \beta_{hh'} x_{ih} x_{ih'} + E_i \quad (1)$$

($i = 1, \dots, n$),

with

- Y_i : simulation response of factor combination i (stochastic variables are shown in capitals);
- β_0 : overall mean response or regression intercept;
- β_h : main or first-order effect of factor h ;
- x_{ih} : value of the standardized factor h in combination i (see Equation (2) below);
- $\beta_{hh'}$: interaction between factors h and h' with $h \neq h'$;
- β_{hh} : quadratic effect of factor h ;
- E_i : fitting error of the regression model for factor combination i ;
- n : number of simulated factor combinations.

First ignore interactions and quadratic effects, for didactic reasons. Then the *relative importance* of a factor is obtained by sorting the absolute values of the main effects β_h , provided the factors are *standardized*. So let the original (non-standardized) factor h be denoted by z_h . In the simulation experiment z_h ranges between a lowest value l_h and an upper value u_h ; that is, the simulation model is not valid outside that range (see the discussion on validation in §3) or in practice that factor can range over that domain only (for example, the number of servers can vary only between one and five). The variation (or spread) of that factor is measured by $a_h = (u_h - l_h)/2$; its location (or mean) by $b_h = (u_h + l_h)/2$. Then the following standardization is appropriate:

$$x_{ih} = (z_{ih} - b_h)/a_h. \quad (2)$$

The classic *fitting algorithm*, which determines β of the regression model in Equation (1), uses the *ordinary least squares* (OLS) criterion. Software for this algorithm is abundant.

If *statistical assumptions* about the fitting error are added, then there are better algorithms. Consider the following assumptions.

It is realistic to assume that the variance of the stochastic fitting error E varies with the input combination of the random simulation model: $\text{var}(E_i) = \sigma_i^2$.

(So Y , the response of the stochastic simulation, has a mean and a variance that both depend on the input.) Then *weighted least squares* (with the standard deviations σ_i as weights) yields unbiased estimators of the factor effects, but with smaller variances than OLS gives.

Common pseudorandom number seeds can be used to simulate different factor combinations. Then GLS gives minimum variance, unbiased estimators. Unfortunately, in practice the variances and covariances of the simulation responses Y are unknown, so they must be estimated. The following equation gives the classic covariance estimator, assuming d independent replications (or simulation runs) per factor combination (so Y_{ig} and $Y_{i'g}$ are correlated, but Y_{ig} and $Y_{i'g}$ are not):

$$\text{cov}(Y_i, Y_{i'}) = \sum_{g=1}^d (Y_{ig} - \bar{Y}_i)(Y_{i'g} - \bar{Y}_{i'})/(d-1). \quad (3)$$

Fortunately, the resulting *estimated GLS* gives good results; see Kleijnen and Van Groenendaal (1992).

Of course, it is necessary to check the fitted regression metamodel: is it an adequate approximation of the underlying simulation model? Therefore the metamodel may be used to predict the outcomes for *new* factor combinations of the simulation model. So replace β in the specified metamodel by the estimate $\hat{\beta}$, and substitute new combinations of x (there are n old combinations). Compare the predictions \hat{y} with the simulation response y .

A refinement is *cross-validation*: do not add new combinations (which require more computer time), but eliminate one old combination (say combination i) and re-estimate the regression model from the remaining $n-1$ combinations. Repeat this elimination for all values of i (with $i = 1, \dots, n$; see Equation (1)). Statistical details are discussed in Kleijnen and Van Groenendaal (1992).

Applications of regression metamodeling will be discussed below (§2.3 through §4).

2.3 Statistical Designs

The preceding subsection (§2.2.2) used regression metamodels to approximate the input/output behavior of simulation models. Such a metamodel has (say) q regression parameters β , which measure the effects of the k factors; for example, q equals $k+1$ if there are no high-order effects; if there are interactions between factors, then q increases with $k(k-1)/2$; and so on.

It is obvious that to get unique, unbiased estimators

of these q effects, it is necessary to simulate at least q factor combinations. Moreover, *which* n combinations to simulate (provided that $n \geq q$), can be determined such that the accuracy of the estimated factor effects is maximized (variance minimized). This is the goal of the statistical theory on DOE (which Fisher started in the 1930s and Taguchi continues today).

2.3.1 Main Effects Only

Consider a first-order polynomial, which is a model with only k main effects, besides the overall mean (see the first two terms in the right-hand side of Equation (1)).

In practice, analysts usually first simulate the 'base' situation, and next they change *one factor at a time*; so, all together they simulate $1 + k$ runs.

However, DOE proves that it is better to use *orthogonal* designs, that is, designs that satisfy

$$\mathbf{x}'\mathbf{x} = n\mathbf{I} \quad (4)$$

with

bold letters: matrices;
 $\mathbf{x} = (x_{ij})$: design matrix with $i = 1, \dots, n$; $j = 0, 1, \dots, k$; $n > k$;
 $x_{i0} = 1$: dummy factor;
 x_{ih} : see text below Equation (1);
 \mathbf{I} : identity matrix (this capital letter does not denote a stochastic variable).

Orthogonal designs give estimators of β that are unbiased and have smaller variances than the estimators resulting from designs that change one factor at a time.

Orthogonal designs are tabulated in many publications. The analysts may also learn how to construct those designs; see Kleijnen (1987). Recently, software has been developed to help the analysts specify these designs; see Ören (1993).

A well-known class of orthogonal designs are 2^{k-p} fractional factorials. An example is a simulation with $k = 7$ factors, which requires $n = 2^{7-4} = 8$ factor combinations (runs). Actually, these 2^{k-p} designs also require 8 runs when $4 \leq k \leq 7$. See Kleijnen (1987).

References to many *simulation applications* of these designs can be found in Kleijnen (1987) and Kleijnen and Van Groenendaal (1992).

In practice, however, it is unknown whether only main effects are important. Therefore orthogonal designs with $n \approx k + 1$ should be used only in optimization (see §4). Moreover these designs are useful as building blocks if interactions are accounted for; see the next sub-subsection (§2.3.2).

2.3.2 Main Effects Biased by Interactions?

It seems prudent to assume that *interactions* between pairs of factors may be important. Then the k main effects can still be estimated without bias caused by these interactions. However, the number of simulated factor combinations must be *doubled*; for example, $k = 7$ requires $n = 2 \times 8 = 16$. These designs also give an indication of the importance of interactions; also see the next sub-subsection (§2.3.3).

Details, including *simulation applications* are presented in Kleijnen (1987) and Kleijnen and Van Groenendaal (1992).

Recent applications include the simulation of a decision support system (DSS) for the investment analysis of gas pipes in Indonesia, and a simulation model for the Amsterdam police; see Van Groenendaal (1994) and Van Meel (1994) respectively.

2.3.3 Factor Interactions

Suppose the analysts wish to estimate the *individual two-factor interactions* β_{hh} ; see Equation (1). There are $k(k-1)/2$ such interactions. Then many more simulation runs are necessary. An example is $k = 7$, which requires $n = 2^{7-1} = 64$ factor combinations (runs). Therefore only small values for k are studied in practice. Kleijnen (1987) gives details, including applications.

Of course, if k is really small (say, $k = 3$), then all 2^k combinations are simulated, so all interactions (not only two-factor interactions) can be estimated. In practice, these *full factorial* designs are sometimes used indeed (but high-order interactions are hard to interpret). See Kleijnen (1987).

2.3.4 Quadratic Effects: Curvature

If the quadratic effects β_{hh} in Equation (1) are to be estimated, then at least k extra runs are needed (since $h = 1, \dots, k$). Moreover, each factor must be simulated for more than two values.

Popular in statistics and in simulation are *central composite designs*. They have five values per factor, and require many runs ($n \gg q$). For example, if there are $k = 2$ factors, then $q = 6$ effects are to be estimated but $n = 9$ factor combinations are simulated. See Kleijnen (1987) and Kleijnen and Van Groenendaal (1992).

Applications are found in the *optimization* of simulation models; see §4.

3 VALIDATION

This paper is confined to the role of sensitivity analysis (§2) in validation; other statistical techniques for validation and verification are discussed in Kleijnen (1995a). Obviously, validation is one of the first questions that must be answered in a simulation study; for didactic reasons, validation is discussed in this part of the paper.

True validation requires that *data* on the real system be available. In practice, the amount of data varies greatly: data on failures of nuclear installations are rare, whereas electronically captured data on computer performance and on supermarket sales are abundant.

If data are available, then many statistical techniques can be applied. For example, simulated and real data on the response, can be compared through the Student statistic for paired observations, assuming the simulation is fed with real-life input data: *trace driven simulation*. A better test uses regression analysis; see Kleijnen, Bettonvil, and Van Groenendaal (1995).

However, if no data are available, then the following type of sensitivity analysis can be used. The clients of the analysts do have *qualitative* knowledge of certain parts of the real system; that is, these clients do know in which direction certain factors affect the response of the corresponding module in the simulation model (also see the discussion on sequential bifurcation in §2.1.1). If the regression metamodel (see §2.2.2) gives an estimated factor effect with the *wrong sign*, this is a strong indication of a wrong simulation model or a wrong computer program.

Applications in ecological and military modeling are given in Kleijnen, Van Ham, and Rotmans (1992) and Kleijnen (1995b) respectively. These applications further show that the validity of a simulation model is restricted to a certain domain of factor combinations, which corresponds with the *experimental frame* in Zeigler (1976), a seminal book on modeling and simulation.

Moreover, the regression metamodel shows which factors are most important. If possible, information on these factors should be collected, for validation purposes.

4 OPTIMIZATION: RSM

There are many mathematical techniques for finding optimal values for the decision variables of nonlinear implicit functions (such as are formulated by simulation models), possibly with stochastic noise. Examples of such techniques are genetic algorithms, simulated annealing, and tabu search. However, this paper is

limited to *RSM*.

First consider four general characteristics of RSM; then some details:

(i) RSM relies on first-order and second-order polynomial regression metamodels, now called *response surfaces*; see §2.2.2.

(ii) It uses the *statistical designs* of §2.3.

(iii) It is augmented with the mathematical (not statistical) technique of *steepest ascent*, to determine in which direction the decision variables should be changed.

(iv) It uses the mathematical technique of *canonical analysis* to analyze the shape of the optimal region: does that region have a unique maximum, a saddle point or a ridge?

More specifically, RSM begins by selecting a *starting point*. Because RSM is a heuristic (no success guaranteed), several starting points may be tried later on, if time permits.

RSM explores the *neighborhood* of that point. The response surface is approximated locally by a first-order polynomial in the decision variables (Taylor series expansion).

The main effects β_h are estimated, using a design with $n \approx k + 1$ (see §2.3.1). Suppose $\hat{\beta}_1 \gg \hat{\beta}_2 > 0$. Then obviously the increase of decision variable 1 (say) z_1 should be larger than that of z_2 . The *steepest ascent path* means $\Delta z_1 / \Delta z_2 = \hat{\beta}_1 / \hat{\beta}_2$ (no standardization; also see next paragraph).

Unfortunately, the steepest ascent technique does not quantify the *step size* along this path. Therefore the analysts may try a specific value for the step size. If that value yields a lower response, then this value should be reduced. Otherwise, one more step is taken. Ultimately, the response must decrease, since the first-order polynomial is only an approximation. Then the procedure is *repeated*: around the best point so far, a new first-order polynomial is estimated, after simulating $n \approx k + 1$ combinations of z_i through z_k . And so on.

In the neighborhood of the top, a hyperplane can *not* be an adequate representation. Cross-validation may be used to detect this lack of fit. Other diagnostic measures are $R^2 \ll 1$ (where R^2 denotes the multiple correlation coefficient), $\text{var}(\hat{\beta}_h) \gg \hat{\beta}_h$, and modern statistics such as PRESS, discussed in Kleijnen (1987).

So when a hyperplane no longer approximates the local input/output behavior well enough, then a second-order polynomial is fitted; see §2.3.4.

Finally, the *optimal* values of z_h are found by straightforward differentiation of the fitted quadratic polynomial. A more sophisticated evaluation is canonical analysis.

Consider the following *case study*. A decision support system (DSS) for production planning in a steel tube factory is simulated and is to be optimized. There are fourteen decision variables, and two response variables (namely, a production and a commercial criterion). Simulation of one combination takes six hours of computer time, so searching for the optimal combination must be performed with care. Details can be found in Kleijnen (1993).

More applications can be found in Hood and Welch (1993), Kleijnen (1987), and Kleijnen and Van Groenendaal (1992).

5 CONCLUSIONS

In the Introduction (§1) the following *questions* were raised:

1. *What if*: what happens if the analysts change parameters, input variables or modules of a simulation model? This question is closely related to sensitivity analysis and optimization.
2. *Validation*: is the simulation model an adequate representation of the corresponding system in the real world?

These questions were answered as follows.

In the initial phase of a simulation it is often necessary to perform *screening*: which factors among the multitude of potential factors are really important? Screening's aim is to reduce the number of really important factors to be further explored in the next phase. The technique of sequential bifurcation is a simple, efficient, and effective screening technique.

Once the important factors are identified, further analysis with fewer assumptions (no known signs) may use *regression analysis*. It generalizes the results of the simulation experiment, since it characterizes the input/output behavior of the simulation model.

Design Of Experiments (DOE) can give good estimators of the main effects, interactions, and quadratic effects that occur in the regression model. These designs require relatively few simulation runs.

Once these factor effects are quantified, they can be used in

- (i) *validation*, especially if there are no data on the input/output of the simulation model or its modules;
- (ii) *optimization* through RSM, which builds on regression analysis and experimental designs.

These statistical techniques have already been *applied* many times in practical simulation studies, in many domains. Hopefully, this survey will stimulate even more analysts to apply these techniques. The goal is to make simulation studies give more general results, in less time.

In the mean time the *research* on statistical techniques adapted to simulation, continues in both Europe and the USA.

REFERENCES

- Bettonvil, B. and J.P.C. Kleijnen. 1994. Identifying the important factors in simulation models with many factors. Tilburg University.
- Glynn, P.W. and D.L. Iglehart. 1989. Importance sampling for stochastic simulation. *Management Science* 35: 1367-1392.
- Ho, Y. and X. Cao. 1991. *Perturbation analysis of discrete event systems*. Dordrecht: Kluwer.
- Hood, S.J. and P.D. Welch. 1993. Response surface methodology and its application in simulation. In *Proceedings of the 1993 Winter Simulation Conference*.
- Kleijnen, J.P.C. 1987. *Statistical tools for simulation practitioners*. New York: Marcel Dekker.
- Kleijnen, J.P.C. 1993. Simulation and optimization in production planning: a case study, *Decision Support Systems* 9: 269-280.
- Kleijnen, J.P.C. 1995a. Verification and validation of simulation models. *European Journal of Operational Research* 82: 145-162.
- Kleijnen, J.P.C. 1995b. Statistical validation of simulation models: a case-study. *European Journal of Operational Research* (in press).
- Kleijnen, J.P.C. 1996. Simulation: sensitivity analysis and optimization through regression analysis and experimental design. In *Proceedings of NATO Advanced Study Institute on Current Issues and Challenges in the Reliability and Maintenance of Complex Systems*, Heidelberg: Springer-Verlag.
- Kleijnen, J.P.C., B. Bettonvil, and W. Van Groenendaal. 1995. Validation of simulation models: regression analysis revisited. Tilburg University.
- Kleijnen, J.P.C., G. Van Ham, and J. Rotmans. 1992. Techniques for sensitivity analysis of simulation models: a case study of the CO2 greenhouse effect. *Simulation* 58: 410-417.
- Kleijnen J.P.C. and W. Van Groenendaal. 1992. *Simulation: a statistical perspective*. Chichester (U.K.): Wiley.
- Ören, T.I. 1993. Three simulation experimentation environments: SIMAD, SIMGEST, and E/SLAM. In *Proceedings of the 1993 European Simulation Symposium*. La Jolla: Society for Computer Simulation.
- Rubinstein, R.Y. and A. Shapiro. 1993. *Discrete event systems: sensitivity analysis and stochastic optimization via the score function method*, New York:

Wiley.

Van Groenendaal, W. 1994. *Investment analysis and DSS for gas transmission on Java*. Tilburg (Netherlands): Tilburg University.

Van Meel, J. 1994. *The dynamics of business engineering*. Delft (Netherlands): Delft University.

Zeigler, B. 1976. *Theory of modelling and simulation*. New York: Wiley Interscience.

AUTHOR BIOGRAPHY

JACK P.C. KLEIJNEN is Professor of Simulation and Information Systems in the Department of Information Systems and Auditing; he is also associated with the Center for Economic Research (CentER). Both the Department and the Center are within the School of Management and Economics of Tilburg University (Katholieke Universiteit Brabant) in Tilburg, Netherlands. He received his Ph.D. in Management Science at Tilburg University. His research interests are in simulation, mathematical statistics, information systems, and logistics. He published six books and more than 130 articles; lectured at numerous conferences throughout Europe, the USA, Turkey, and Israel; he was a consultant for various organizations; and is a member of several editorial boards. He spent some years in the USA, at different universities and companies. He was awarded a number of fellowships, both nationally and internationally.