Proceedings of the 1995 Winter Simulation Conference
ed. C. Alexopoulos, K. Kang, W. R. Lilegdon, and D. Goldsman

WORKING SMARTER WHEN DEVELOPING LINEAR SIMULATION METAMODELS

Michael K. Taylor

Air Force Studies & Analysis Agency
Force Applications Directorate
Washington, DC 20330-1570, U.S.A.

Paul F. Auclair
Edward F. Mykytka

Department of Operational Sciences
Air Force Institute of Technology / ENS
2950 P Street; Building 640
Wright-Patterson AFB, OH 45433-7765, U.S.A.

ABSTRACT

In this paper, we propose that proper metamodel specification is far more important than the level of simulation effort used in developing a metamodel to estimate the expected response of a given system or simulation. In particular, M/M/k queues with various configurations of arrival rate, service rate, and number of servers were simulated using different levels of simulation effort. The average queue length for each configuration was computed, and then two different metamodels were fit to the simulation data. Analysis of the residuals from the fitted metamodels indicates that metamodel specification has a significant effect on the statistical quality of the estimated expected queue length while the level of simulation effort used in fitting the metamodel has virtually no such effect.

1 INTRODUCTION

Computer simulation is often used to predict the future behavior of a real system as part of an overall experimental modeling process (Shannon 1992, p. 65). This prediction often takes the form of estimates of the expected outputs of the system in response to a given set of inputs. In general, the goodness or statistical precision of these estimates is an increasing function of the amount of "simulation work" performed. That is, their precision can be improved by

1. increasing the number of replications performed, and/or
2. increasing the length of each replication, when the objective is to estimate the long-run (or steady state) expected output of the system.

Thus, "more work is better" and most simulation practitioners would likely prefer to conduct the largest number of replications with the longest run lengths that resources would permit.

Unfortunately, it can be costly to run the simulation model each time a new set of inputs is to be examined. In such cases, Kleijnen (1987) suggests the use of a "metamodel" as a surrogate for the simulation model. A metamodel is generally a mathematical model that attempts to empirically express the output from a simulation as a function of its inputs. Its objective is to "effectively relate the output data of a simulation model to the model's input to aid in the purpose for which the simulation model was developed" (Sargent 1992, p. 888), namely estimating the expected output of the system. Once developed, a metamodel can be used in lieu of the simulation to estimate the desired system performance characteristic for each new set of inputs.

The use of a metamodel as a surrogate for a simulation model results in savings in terms of both computer run time and analysis time. These savings are obtained, however, at the expense of precision. A metamodel, being an empirical representation of previously observed data from the simulation, only provides *estimates* of the expected *simulation* output.

The ability of a metamodel to provide precise estimates of the expected simulation output and, in turn, of the expected output of the real system, would appear to be dependent on how well the metamodel is specified , i.e., on how well it approximates the "true" input-output mapping of the simulation model. A "good" metamodel also would be parsimonious, providing acceptable estimates of the simulation output while containing as few terms or parameters as possible.

The quality of a metamodel would thus appear to depend on both the statistical quality and the amount of the data initially made available for its fitting which, in turn, would depend on the amount of "simulation work" performed in generating the data. Our initial conjecture was twofold:

1. "more" simulation work would produce "better" metamodels, and
2. "better" metamodels would produce "better" estimates of the expected system output.

2 A SIMULATION EXPERIMENT

The purpose of this research was to examine the effects of simulation work and metamodel specification on the statistical quality of the estimates of the mean output of the system provided by a metamodel. Toward this end, a simulation model was developed for an M/M/k queuing system with the objective of estimating the expected length of the waiting line. (This system was chosen since this performance characteristic could be computed analytically using standard queuing results.) The goal was to use this simulation model to develop a database upon which metamodels could be fit and subsequently be used to estimate the expected length of the waiting line for queuing systems within the design region specified by Table 1.

Table 1: Queuing System Configurations

Configuration	Arrival Rate	Service Rate	Number of Servers	Utilization Rate
1	1.0	1.0	2	0.5
2 *	1.0	1.0	4	0.25
3 *	1.0	1.25	2	0.4
4	1.0	1.25	4	0.2
5 *	1.5	1.0	2	0.75
6	1.5	1.0	4	0.375
7	1.5	1.25	2	0.6
8 *	1.5	1.25	4	0.3

* Configurations used in half-fraction design

Table 1 displays the eight different queuing configurations that were used to develop the metamodel database. These configurations differ according to the values of three input parameters—the mean arrival rate, the mean service rate, and the number of servers. The levels of these three parameters were chosen to provide a reasonably broad range of system utilization rates, as shown in the table.

This design was used as the basis for an experiment in which two different metamodels were fit to the output from the simulation model using different levels of simulation work. The effects of simulation work and metamodel specification were then assessed by comparing the estimates of average queue length produced by the metamodels with the expected queue length computed analytically.

2.1 Levels of Simulation Work

Table 2 shows how the levels of simulation work were varied by setting the simulation run length and the number of replications. In this table, the value shown for the simulation "work" is the product of the simulation length (measured in terms of the total time simulated) and the number of simulation replications. This measure of simulation work is somewhat arbitrary. Alternate measures are certainly possible, including one in which simulation length is measured in terms of the number of customers served (which may more accurately reflect processing time). Our particular measure was chosen for its simplicity; we do not expect alternate measures to substantially change the inferences made from our experiment.

Table 2: Final Simulation Work Combinations

Case	Number of Replications	Simulation Length	Simulation Work
A	5	2,500	12,500
A1	5	1,500	7,500
A2	5	1,250	6,250
A3	5	1,000	5,000
A4	5	750	3,750
A5	5	500	2,500
A6	5	250	1,250
C	5	10,000	50,000
C1	5	20,000	100,000
E	10	5,000	50,000
G	20	2,500	50,000
I	20	10,000	200,000

The levels chosen for the simulation length were based on Nelson's suggestion that a simulation length of 20 times the length of the initial transient be used (Nelson 1992, p. 130). Preliminary analysis indicated that a *conservative* estimate for the initial transient period was 500 time units. Thus, the "acceptable" run length was set to 10,000 time units. Further, since increasing the simulation run length would only serve to improve an already acceptable simulation estimate, this research focused on run lengths substantially *less* than 10,000 time units.

The levels chosen for the number of simulation replications were based on the common acceptance of 30 as an effective sample size by most statistical practitioners (Mendenhall, Wackerly, and Scheaffer 1990, p. 319). To examine the effect of simulation work on the statistical quality of the metamodel estimates, this research considered sample sizes substantially *less* than 30.

The twelve combinations of run length and number of replications shown in Table 2 were selected after some preliminary analysis associated with an initial 2-factor, 3-level design in which the levels of run length were set to 2,500, 5,000, and 10,000 time units while those for the number of replications were set to 5, 10, and 20. Of these, five were retained for our experiment and are labeled as cases A, C, E, G, and I in Table 2. The five cases were the four combinations of the low and high factor levels and the combination of the "middle" factor levels–the corners and "center" point of the initial design.

Preliminary analysis suggested that it would be worthwhile to investigate seven additional cases. Six of these (labeled as cases A1 through A6) were obtained by decreasing the run length in case A (which had the smallest amount of work in the preliminary design) in order to study the effect of even less simulation work. In fact, cases A5 and A6 have run lengths that are less than or equal to our conservative estimate of the length of the initial transient period. The seventh additional case was obtained by increasing the run length in case C and is labeled as C1. It provides a run length well beyond that suggested by Nelson and represents an extreme condition intended to assess the validity of the original upper bound on run length in the experiment.

2.2 Metamodel Data Bases

Two different data bases upon which metamodels could be fit were generated for each case of simulation work shown in Table 2. The first of these consists of the output from the full-factorial experiment described by Table 1 while the second consists of the output from the half-fraction whose configurations are highlighted with an asterisk in Table 1. The block generator for the half-fraction is the three-factor interaction. The choice of the particular half-fraction indicated was somewhat arbitrary–its only distinguishing aspect is that its range of system utilization rates is broader than that of the other half-fraction with the same block generator. The output from the two designs will be referred to, respectively, as the Full and Fractional data bases.

2.3 Metamodels Selected

The metamodels used in this paper were formulated by Friedman and Friedman (1985), who validated the use of the following multiplicative metamodel for estimating the average queue length of M/M/k queues:

$$L_q = e^{\beta_0} Arr^{\beta_1} Svc^{\beta_2} Num^{\beta_3},$$

where L_q, *Arr*, *Svc*, and *Num* respectively denote the expected queue length, the mean arrival rate, the mean service rate, and the number of servers. This multiplicative metamodel is used in this research as an example of a "well-specified" metamodel–one that can be expected to provide an adequate fit of the simulation data. Additionally, Friedman and Friedman's linear metamodel, which has the form

$$L_q = \beta_0 + \beta_1 Arr + \beta_2 Svc + \beta_3 Num ,$$

is used in this research as an example of a "poorly specified" metamodel–one that would *not* be generally expected to provide an adequate fit of the simulation data. Note that the two metamodels are "equally parsimonious" in the sense that each has the same number of parameters (four).

2.4 The Experiment

To illustrate the experimental process, we next show how metamodels were developed for case A from Table 2. First, for each of the eight configurations displayed in Table 1, five replications of the simulation (each of length 2,500) were performed. The average queue lengths observed in these 40 experiments are shown in Table 3.

Table 3: Average Queue Lengths, Case A

	Replication				
Config.	1	2	3	4	5
1	0.361	0.356	0.328	0.366	0.376
2	0.003	0.008	0.005	0.008	0.005
3	0.186	0.155	0.157	0.139	0.146
4	0.005	0.003	0.003	0.002	0.004
5	2.037	1.437	1.786	2.385	1.816
6	0.038	0.046	0.051	0.042	0.040
7	0.637	0.813	0.657	0.815	0.687
8	0.025	0.010	0.017	0.020	0.015

For each configuration, the average queue length across the five replications was computed. This was subtracted from the expected queue length for that configuration (which was computed analytically) to form a residual. The results for case A are depicted in Table 4.

The simulation output was then used to fit both the multiplicative and the linear metamodels to the data developed within each case for each of the two databases. The regression coefficients for the predictive metamodels were estimated using least squares regression. The least squares fit for the multiplicative metamodel was obtained in reference to the transfomed model:

$$\ln(L_q) = \beta_0 + \beta_1 \ln(Arr) + \beta_2 \ln(Svc) + \beta_3 \ln(Num).$$

Table 4: Summary Statistics, Case A

Config.	Analytical Solution	Average Length	Residual
1	0.3333	0.3574	-0.0241
2	0.0068	0.0058	0.0010
3	0.1524	0.1566	-0.0042
4	0.0024	0.0034	-0.0010
5	1.9286	1.8922	0.0364
6	0.0448	0.0434	0.0014
7	0.6750	0.7218	-0.0468
8	0.0159	0.0174	-0.0015

In both cases, the metamodels were fit using the simulation data obtained from the full-factorial design depicted in Table 1.

Once these metamodels were fit, they were used to estimate the expected queue length for each of the eight queuing configurations. A summary of the estimated expected queue lengths for each configuration, metamodel, and database is provided in Table 5.

Table 5: Estimated Expected Queue Lengths, Case A

Cfg	Metamodel					
	Full Log	Full Lin-N	Full Lin-Z	Frac Log	Frac Lin-N	Frac Lin-Z
1	0.329	0.685	0.685	0.303	1.019	1.019
2	0.007	-.078	**0**	0.005	0.006	0.006
3	0.146	0.337	0.337	0.298	0.157	0.157
4	0.003	-.426	**0**	0.005	-.865	**0**
5	1.848	1.224	1.224	1.834	1.892	1.892
6	0.038	0.461	0.461	0.033	0.879	0.879
7	0.822	0.876	0.876	1.803	1.030	1.030
8	0.017	0.113	0.113	0.033	0.017	0.017

In Table 5, the columns denoted "Full Log" and "Frac Log" contain the estimates of expected queue length computed using the multiplicative metamodel as fit over the full-factorial and the fractional-factorial databases, respectively. Similarly, the columns denoted "Full Lin-N" and "Frac Lin-N" contain the estimates of expected queue length computed using the linear metamodel fit over each of the same two databases. It can be observed, however, that the linear metamodel sometimes produces *negative* estimates of the expected queue length (see, for instance, configuration 4). For this reason, the "Full Lin-N" and "Frac Lin-N" sets of estimates are referred as the "linear-negative" sets of estimates as fit over the full and fractional databases, respectively. Further, since negative estimates would be inappropriate or unacceptable in many practical situations, we created a second set of linear metamodel estimates wherein any negative values are reset to zero. These estimates are provided in the columns headed "Full Lin-Z" and "Frac Lin-Z" and are referred to as the "linear-zero" sets of estimates. Thus, in the linear-negative sets, any negative estimates of average queue length are retained while, in the linear-zero sets, these are reset to zero. (No preference is given to either of these sets in the following analysis.)

Next, for each metamodel and configuration, residuals were calculated as the differences between the expected queue lengths (as computed analytically) and those estimated using the corresponding metamodel. These residuals are shown in Table 6.

Table 6: Residuals, Case A

Cfg	Metamodel					
	Full Log	Full Lin-N	Full Lin-Z	Frac Log	Frac Lin-N	Frac Lin-Z
1	0.329	0.685	0.685	0.303	1.019	1.019
2	0.007	-.078	0.000	0.005	0.006	0.006
3	0.146	0.337	0.337	0.298	0.157	0.157
4	0.003	-.426	0.000	0.005	-.865	0.000
5	1.848	1.224	1.224	1.834	1.892	1.892
6	0.038	0.461	0.461	0.033	0.879	0.879
7	0.822	0.876	0.876	1.803	1.030	1.030
8	0.017	0.113	0.113	0.033	0.017	0.017

3 RESULTS & ANALYSIS

The preceding process was repeated for each of the remaining 11 cases of simulation work shown in Table 2. Summary statistics describing the residuals across the eight configurations within each case are graphically presented in Figures 1, 2, and 3. These figures, respectively, plot the sample means, standard errors, and root mean squared errors of the residuals across all configurations within each case and metamodel type. As a basis for comparison, the figures also include these same statistics for the estimated queue length obtained directly from the simulation output for each corresponding metamodel. These figures provide a subjective means of evaluating the overall quality of the estimates provided by each of the different types of metamodel as a function of the amount of simulation work performed in developing that metamodel.

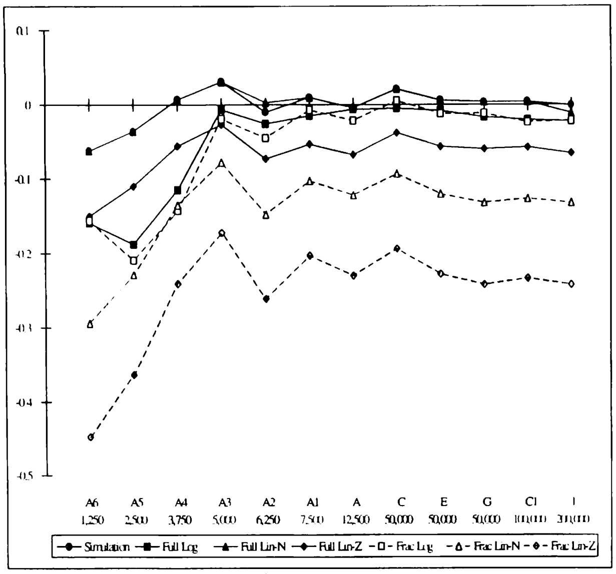

Figure 1: Mean of Residuals by Simulation Work Case

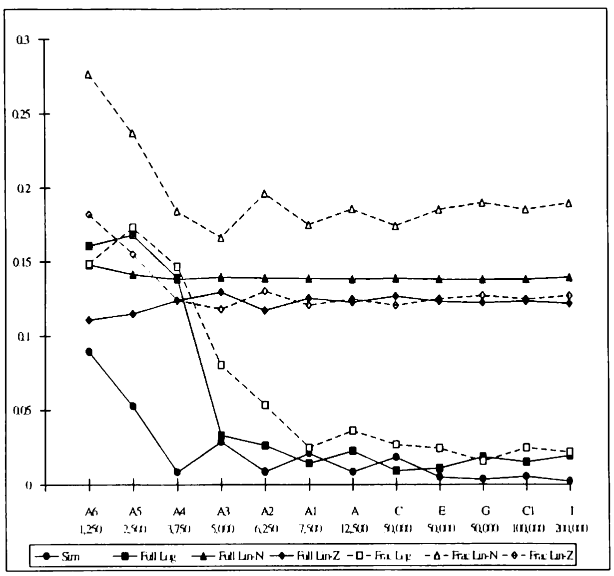

Figure 2: Standard Error of Residuals by Simulation Work Case

3.1 Means of the Residuals

In particular, the sample means of the residuals for each case of simulation work are depicted in Figure 1. The most striking aspect of this figure appear to be that, once a minimum amount of simulation work is performed, the means of the residuals do not substantially change as the amount of simulation work is increased. For example, the

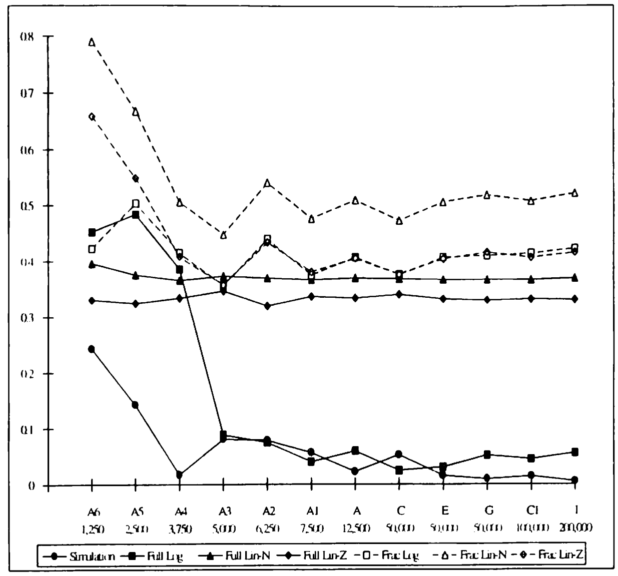

Figure 3: Root Mean Squared Error of Residuals by Simulation Work Case

means of the residuals from the Full Linear-Zero case vary around an average of about -0.05 over all cases of simulation work in excess of 2,500 units (i.e., over all cases except A6 and A5, which have means of -0.150 and -0.110, respectively). Thus, once a minimum amount of simulation work is performed, increasing the amount of work does not appear to substantially affect the average quality of the estimates provided by the metamodels.

It is also apparent that, for a specific case of simulation work, the means of the residuals do differ substantially between metamodels, as indicated by the fact that the "curves" traced out by each metamodel tend to be parallel to one another. This difference seems to suggest that the specification of a metamodel *does* appear to affect the average quality of the estimates provided by that metamodel.

Finally, note that for all but the smallest three cases of simulation work, both multiplicative metamodels (i.e., those developed using either the full- or fractional-factorial databases) and the linear-negative metamodel appear to produce estimates whose average residuals are nearly identical to those produced by the simulation model itself. This would suggest that each of these metamodels is well-specified and could be used as a viable surrogate for the simulation model.

3.2 Standard Errors of the Residuals

The standard errors of the residuals for each case are graphically presented in Figure 2. It should be immediately apparent that the standard errors tend to initially decline as the amount of simulation work is increased but that these quickly tend to level off. Thus, it again appears that, once a minimum amount of simulation work is performed, increasing the amount of work does not appear to substantially affect the average quality of the estimates provided by the metamodels. Further, once this threshold of simulation work is reached, the metamodels tend to reside in one of two groups: those with standard errors of less than 0.05 and those with standard errors greater than 0.10. It is, perhaps, no surprise that the metamodels with the smallest standard errors have a multiplicative form, demonstrating that the multiplicative metamodels are indeed "well-specified." This price paid for mis-specifying the metamodel is seen in the larger standard errors associated with the linear metamodels.

It is interesting to note additionally that, while the linear-negative metamodel has very small residuals on the average, it also has relative large standard errors. This demonstrates that the average residual, by itself, may not be a good measure of the quality of the estimates produced by a metamodel. Finally, it is also interesting to observe that the two multiplicative metamodels perform similarly with respect to both their average residuals and their standard errors, with the metamodel based on the full factorial database consistently performing at least as well as that based on the fractional factorial database.

3.3 Root Mean Squared Errors

In order to try to capture both of the preceding measures of the quality (the mean and standard error) in one statistic, we also computed the root mean squared errors of the residuals. The root mean squared errors are depicted in Figure 3. As perhaps should be expected, the root mean squared errors initially decline as the amount of simulation work is increased but level off rather quickly. This again is evidence that, once a minimum amount of simulation work is performed, increasing the amount of work does not substantially affect the average quality of the estimates provided by the metamodels. Interestingly, the effects of small amounts of simulation work on quality appear to be somewhat more pronounced here than in the previous cases, demonstrating the importance of utilizing some minimum level of simulation work. It is notable, however, that this minimum appears to be *much* smaller than we would have obtained had we decided to run our simulation using a run length of 10,000 time units, as suggested by the application of our conservative estimate of the warm-up period to Nelson's criterion. (If five replications were performed, this would correspond to 50,000 units of simulation work, as depicted by our case C.)

As in Figure 2, once this threshold of simulation work is reached, the curves in Figure 3 suggest that the metamodels tend to cluster in one of two groups: those with root mean squared errors of less than 0.10 and those with root mean squared errors greater than 0.30. What is striking in this case is that only one metamodel–the multiplicative metamodel based on the full-factorial database–falls into this first group. Further, its root mean squared errors are similar to those produced by estimates derived from the simulation itself, suggesting that this metamodel is in a class by itself as a surrogate for the simulation. The fact that the multiplicative metamodel based on the fractional database now falls in the latter group demonstrates, perhaps, the price paid for using a smaller database. Further, since its root mean squared errors tend to exceed those associated with the linear metamodels based on the full-factorial database, it appears that the reduction in quality associated with model mis-specification is comparable to the reduction in quality associated with using a smaller database.

3.4 Analyses of Variance

Since the preceding graphical analysis is somewhat subjective, one-way analyses of variance (ANOVA's) were performed to determine if the statistical quality of the estimates produced by the metamodels significantly differ according to either

(i) the type of metamodel used,
(ii) the database used, or
(iii) the amount of simulation work performed.

For the purpose of these analyses, we measure statistical quality in terms of the root mean squared error.

The first ANOVA was performed to determine if the root mean squared errors for the six types of metamodels (Full-Log, Fractional-Log, Full-Linear-Negative, Full-Linear-Zero, Fractional-Linear-Negative, and Fractional-Linear-Zero) are not significantly different from each other. The ANOVA table depicted in Table 7 shows that there is quite clearly a difference between metamodel types and establishes that the statistical quality of the estimates produced by a metamodel does indeed depend on proper metamodel specification.

Table 7: ANOVA Table–Types of Metamodels

Source of Variation	SS	df	MS	F	P-value	F crit
Between Groups	1.008	5	0.201	24.33	0.000	2.354
Within Groups	0.547	66	0.008			
Total	1.554	71				

The second ANOVA was performed to determine if the root mean squared errors significantly differ according to which database–full or fractional–was used in developing a metamodel. The corresponding AVOVA table is presented in Table 8 and shows that there is a significant difference attributable to the database used.

Table 8: ANOVA Table–Metamodel Databases

Source of Variation	SS	df	MS	F	P-value	F crit
Between Groups:	0.564	1	0.564	39.85	0.000	3.978
Within Groups:	0.991	70	0.014			
Total:	1.554	71				

The third ANOVA was performed to determine if the root mean squared errors significantly differed according to the amount of simulation work performed. As seen from Table 2, there are 10 different levels of work used in our experiment–1,250, 2,500, 3,750, 5,000, 6,250, 7,500, 12,500, 50,000, 100,000, and 200,000. The corresponding ANOVA table is presented as Table 9. These results in this case are weak–the p-value of approximately 0.25 indicates that we would fail to reject the hypothesis that there are no significant differences as a function of the amount of work at a level of significance of 0.25 or less.

Table 9: ANOVA Table–Levels of Simulation Work

Source of Variation	SS	df	MS	F	P-value	F crit
Between Groups:	0.248	9	0.028	1.311	0.250	2.035
Within Groups:	1.306	62	0.021			
Total:	1.554	71				

Since the preceeding conclusion is somewhat weak, and since our graphical analysis suggested that the average quality of the estimates provided by the metamodels does not appear to be substantially affected by increasing the amount of simulation work beyond a minimum level, we next repeated the third ANOVA after deleting the three smallest levels of work. The corresponding ANOVA table is presented as Table 10. Here the results are quite strong–the p-value of 0.9996 indicates quite convincingly that there are no significant differences in the root mean squared errors of the residuals for the different levels of simulation work. We thus conclude that, overall, the amount of simulation work does *not* affect the statistical quality of the estimates produced by a metamodel, with the caveat that this could generally be expected to apply only after a minimum amount of simulation work is performed.

Table 10: ANOVA Table–Seven Highest Levels of Simulation Work

Source of Variation	SS	df	MS	F	P-value	F crit
Between Groups	0.006	6	0.001	0.043	0.999	2.299
Within Groups	1.056	47	0.023			
Total	1.062	53				

4 CONCLUSIONS

The graphical analyses and the ANOVA's support the following two conclusions. First, *the amount of simulation work (beyond a reasonable minimum) has no significant effect on the statistical quality of the estimates produced by a metamodel* while, second, *metamodel specification does have a significant effect on that quality.* These conclusions suggest that metamodel specification is more important than the amount of simulation work performed in fitting a metamodel. Thus, when using a metamodel as a surrogate for a computer simulation, it appears to be more beneficial to expend effort toward developing a better metamodel than toward performing more simulation work in developing a database for fitting that metamodel. In other words, a simulation practitioner is advised to *work smarter, not harder.*

REFERENCES

Friedman, L. W. and H. H. Friedman. 1985. Validating the simulation metamodel: Some practical approaches, *Simulation*, 45: 144-146.

Kleijnen, J. P. C. 1991. *Statistical tools for simulation practitioners.* New York: Marcel Dekker, Inc.

Law, A. M. and W. D. Kelton. 1991. *Simulation modeling & analysis.* 2nd ed. New York: McGraw-Hill, Inc.

Mendenhall, W., D. D. Wackerly, and R. L. Scheaffer. 1990. *Mathematical statistics with applications.* 4th ed. Boston: PWS-Kent.

Nelson, B. L. Designing efficient simulation experiments. 1992. In *Proceedings of the 1992 Winter Simulation Conference*, ed. J. J. Swain, D. Goldsman, R. C. Crain, and J. R. Wilson, 126-132. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

Shannon, R. E. Introduction to simulation. 1992. In *Proceedings of the 1992 Winter Simulation Conference*, ed. J. J. Swain, D. Goldsman, R. C. Crain, and J. R. Wilson, 65-73. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

AUTHOR BIOGRAPHIES

MICHAEL K. TAYLOR is a Theater Command and Control Analyst assigned to the Air Force Studies and Analyses Agency at the Pentagon, Washington DC. He received a B.S. in Mathematics from the University of Alabama at Birmingham in 1986, and an M.S. in Operations Research from the Air Force Institute of Technology in 1994. His research interests include Defense Systems Simulation, Response Surface Methodology, Design of Experiments, Metamodels, and Statistical Aspects of Simulation.

PAUL F. AUCLAIR serves as the Head of the Department of Operational Sciences in the Graduate School of Engineering at the Air Force Institute of Technology. He received a B.S. from the Air Force Academy in 1976; an M.S.S.M. from the University of Southern California in 1979; an M.S. in Operations Research from the Air Force Institute of Technology in 1982; an M.S.S. from the Georgia Institute of Technology in 1989; and a Ph.D., Industrial Engineering, from the Georgia Institute of Technology in 1993. His research interests include Strategic Missile Operations Analysis, Operational Test and Evaluation, CEP Estimation Techniques, Design of Experiments, and Response Surface Methodology.

EDWARD F. MYKYTKA is an Associate Professor of Operations Research and serves as the Deputy Head of the Department of Operational Sciences in the Graduate School of Engineering at the Air Force Institute of Technology. He received a B.S. in Mathematics from the University of Dayton in 1976; an M.S. in Industrial Engineering from the University of Iowa in 1978; and a Ph.D. in Systems Engineering from the University of Arizona in 1983. His research interests include discrete-event system simulation (particularly random variate modeling and generation), quality improvement and statistical process control, stochastic modeling, and applied statistics.