Proceedings of the 1999 Winter Simulation Conference
P. A. Farrington, H. B. Nembhard, D. T. Sturrock, and G. W. Evans, eds.

SENSITIVITY ANALYSIS IN RANKING AND SELECTION FOR MULTIPLE PERFORMANCE MEASURES

Douglas J. Morrice

MSIS Department
The University of Texas
at Austin
Austin, TX
78712-1175 U.S.A.

John Butler

Department of
Accounting and MIS
The Ohio State University
Columbus, OH
43210 U.S.A.

Peter Mullarkey

Maxager Technology
103 Copperleaf Road
Austin, TX
78734 U.S.A.

Srinagesh Gavirneni

Schlumberger
8311 North RR 620
Austin, TX
78726 U.S.A.

ABSTRACT

In this paper, we conduct sensitivity analysis on a ranking and selection procedure for making multiple comparisons of systems that have multiple performance measures. The procedure combines multiple attribute utility theory with ranking and selection to select the best configuration from a set of K configurations using the indifference zone approach. Specifically, we consider sensitivity analysis on the weights generated by the multiple attribute utility assessment procedure. We demonstrate our analysis on a simulation model of a large project that has six performance measures.

1 INTRODUCTION

In recent work, Morrice et al. (1997) present a simulation model of a project that contains multiple input parameters and multiple performance measures. Morrice et al. (1998, 1999) develop a procedure that combines simulation, statistical ranking and selection (R&S), and Multiple Attribute Utility (MAU) theory to select the best project configuration over K (>1) possible configurations. The K configurations are constructed from different settings of the input parameters.

In this paper, we conduct sensitivity analysis on the weights assessed by the MAU procedure. In particular, we perform a one-at-a-time sensitivity analysis on each weight while holding the ratios amongst the other weights constant. We demonstrate the impact of this sensitivity analysis on the indifference zone, the ranking of the configurations, and the total number of simulation runs required. Sensitivity analysis on the MAU weights is an important exercise because the weights represent a quantification of qualitative decision-maker preferences. Thus, their exact values are difficult to know or even assess with certainty.

The remainder of the paper is organized in the following manner. Section 2 describes the project example from Morrice et al. (1997) that will be used throughout the paper. Section 3 describes the methodology developed in Morrice et al. (1998, 1999). Section 4 contains the sensitivity analysis results. Section 5 contains some concluding remarks and discusses future research directions.

2 EXAMPLE

We use the methodology developed in this paper to analyze the simulation model of the project described in Morrice et al. (1997). The simulation models a large outdoor operation called a signal quality survey. Signal quality surveys are conducted over large geographical areas (tens to hundreds of square kilometers). They are projects taking anywhere from a few days to a few years with the number of personnel ranging from 20 to 1000 people, requiring capital equipment valued in the tens of millions of dollars, and generating survey revenues ranging from hundreds of thousands to hundreds of millions of dollars. The simulation model was designed to support bidding, planning, and conducting these large, complicated, and expensive projects in a profitable manner.

The execution of a signal quality survey requires the coordination of five types of crews (see Figure 1). Briefly, the signal crew sends signals from several geographic locations that are recorded by the recording crew. Since the signal crew and recording crew work in such a synchronized fashion, we will consolidate them into one object in this analysis. The layout crew places receiving (or monitoring) equipment at several geographic locations so that the recording crew can receive signals sent by the source crew. The transport crew brings the layout crew receiving equipment. The packing crew prepares receiving equipment for the transport crew that is no longer required on a particular part of a survey for receiving signals sent by the signal crew.

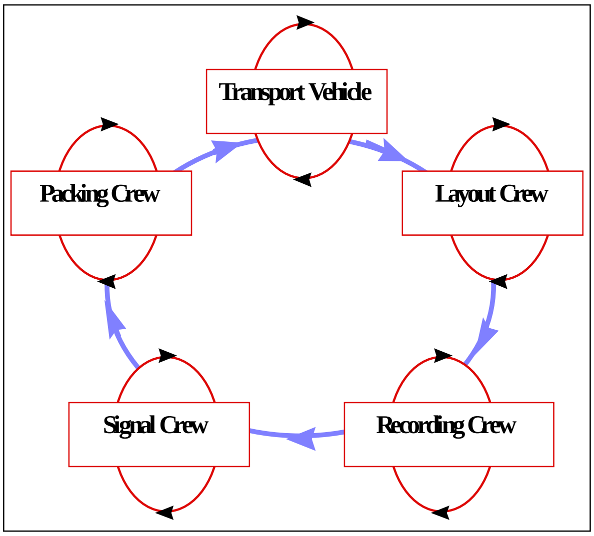

Figure 1: Crews in a Signal Quality Survey

Performance measures on this project include percent utilization for all crew types, project duration, and cost. We will model four project configurations differentiated by the number of source crews and the amount of receiving equipment available. Resource decisions along these two dimensions are considered the most important on a signal quality survey. More specifically, the configurations will be numbered as follows: a single source crew with 1100 units of equipment (1), a single source crew with 1300 units of equipment (2), two source crews with 1100 units of equipment (3), and two source crews with 1300 units of equipment (4). Each configuration will be evaluated on the multiple performance measures and the best will be selected.

3 METHODOLOGY

This section summarizes the methodology developed in Morrice et al. (1998, 1999). It is an R&S technique using the indifference zone approach for multiple performance measures. Assume that there are $K \geq 2$ project configurations and $n \geq 1$ performance measures. For $1 \leq k \leq K$, let $\boldsymbol{X}_k = (X_{k1}, X_{k2},\ldots, X_{kn})$ denote a vector of random variables representing the performance measures for configuration k. Let $E[u(\boldsymbol{X}_k)]$ denote the expected utility (unknown) for configuration k, where $u(\boldsymbol{X}_k)$ is a MAU function of the form:

$$u(\boldsymbol{X}_k) = \sum_{i=1}^{n} w_i u_i(X_{ki}) \qquad (1)$$

The function $u_i(\cdot)$ is a single attribute utility function over measure i that is scaled from 0 to 1, w_i is the weight for measure i and $\sum_{i=1}^{n} w_i = 1$. For more details on MAU theory, see Clemen (1991).

Let

$$E[u(\boldsymbol{X}_{[1]})] \leq E[u(\boldsymbol{X}_{[2]})] \leq \ldots \leq E[u(\boldsymbol{X}_{[K]})]$$

denote the ordered expected utility values. The goal is to select the project configuration with the largest expected utility $E[u(\boldsymbol{X}_{[K]})]$. If the R&S procedure achieves this goal a "correct selection" (CS) is made. The R&S procedure is designed to satisfy the following probability requirement:

$$P\{CS\} \geq P^* \text{ whenever } E[u(\boldsymbol{X}_{[K]})] - E[u(\boldsymbol{X}_{[K-1]})] \geq \delta^*$$

where $(1/K) < P^* < 1$ and $0 < \delta^* < 1$.

To conduct the analysis, a simulation model generates M (≥ 1) replicates for each project configuration. On each replication for each configuration, the multiple attribute utility function in (1) is evaluated using the realization of $\boldsymbol{X}_k$. If the utility function realizations for a given configuration are not normally distributed then multiple replicates are conducted, over which the realizations are averaged. Then multiple replicates are made of the averages in order to produce approximately normal data for the R&S procedure. Goldsman et al. (1991) refer to this last step as making macroreplications. In our analysis, we used the two-stage indifference zone procedure for R&S due to Rinott (1978).

When R&S is based on expected utilities, the selection of δ^* can be challenging because δ^* has no direct physical meaning on the utility scale. To address this problem, Morrice et al. (1999) develop the following procedure:

1. Pick measure j, $j \in \{1,2, \ldots, n\}$ as the standard measure
2. Exchange utility in measure i, $1 \leq i \leq n$,, $i \neq j$ for configuration k, $1 \leq k \leq K$ via:

   $$u_j(X'_{kj}) = u_j(X_{kj}) + \sum_{i \neq j} \frac{w_i}{w_j}[u_i(X_{ki}) - c_i]$$

   where

   - X_{ki} is the original level of performance measure i for configuration k
   - X_{kj}' is the hypothetical level of measure j for configuration k after performing the utility exchange
   - c_i is the common level for measure i, $i \neq j$, in utility, i.e., between 0 and 1

3. Select δ_j^* for the standard measure by inverting $u_j(\cdot)$.

Step 1 selects a single standard performance measure. Step 2 performs a utility exchange transforming all other utilities to a common level and transforming them to the standard measure. Once $u_j(\cdot)$ has been inverted in step 3, an indifference zone is selected on the scale of the original performance measure scale. This indifference zone is then mapped to δ_j^* via $u_j(\cdot)$.

The following two propositions from Morrice et al. (1999), stated without proof, show the relationship between the variance of $u_j(X'_{kj})$ and $u(\boldsymbol{X}_k)$, and the relationship between δ_j^* and δ^*. Proposition 2 provides necessary and sufficient conditions for defining R&S procedures on $u_j(X'_{kj})$ and $u(\boldsymbol{X}_k)$ that are completely equivalent.

Proposition 1: The following relationship holds for the variances:

$$\mathrm{Var}(u_j(X'_{kj})) = \frac{\mathrm{Var}(u(\boldsymbol{X}_k))}{w_j^2}$$

for measure *j*, $j \in \{1,2,\ldots,n\}$ and configuration *k*, $k=1,2,\ldots, K$.

Proposition 2: $\mathrm{E}[u(\boldsymbol{X}_{[\kappa]})]-\mathrm{E}[u(\boldsymbol{X}_{[\kappa-1]})] \geq \delta^*$ iff

$$\mathrm{E}[u_j(X'_{[K]j})] - \mathrm{E}[u_j(X'_{[K-1]j})] \geq \frac{\delta^*}{w_j}$$

for measure *j*, $j \in \{1,2,\ldots,n\}$ and configuration *k*, $k=1,2,\ldots, K$.

Once δ_j^* has been selected, Proposition 2 can also be used for computing the implied δ_i^*, for $1 \leq i \leq n$, $i \neq j$ using the weights assessed from the MAU procedure. By inverting $u_i(\cdot)$, for all *i*, $i=1,2,\ldots, n$, one can construct indifference-zone, preference-zone diagrams (Bechhofer et al. 1995, page 178) for the original performance measures implied by δ_j^* and the weights. Indifference-zone, preference-zone diagrams are helpful to check for consistency of the decision-makers preferences expressed in the weights and single attribute utility curves. They can also be used for conducting sensitivity analysis on the weights, as we will demonstrate in the next section.

4 SENSITIVITY ANALYSIS

To demonstrate sensitivity analysis on the weights, we use the example described in Section 2. This example is described in more detail in Morrice (1999). All performance measures have an exponential utility function of the form:

$$U(X) = A - Be^{(X/RT)} \tag{2}$$

(Clemen 1991, page 379). In particular, project cost has the utility function

$$1.064 - (0.0195)\, e^{(X/50000)}, \tag{3}$$

where X lies in the range of 60 to 200 thousand US dollars. The project duration has as its utility function

$$1.021 - (0.00106)\, e^{(X/80)}$$

where X is between 240 and 550 hours. We use the following utility functions for all the crew ulitizations:

$$1 - (2.061\mathrm{E-}09) e^{(X/0.05)} \tag{4}$$

where X is an element of [0,1].

The MAU function was constructed from a weighted sum of the six single attribute utility functions. Weights were assessed as follows: cost (0.4), job duration (0.2), and worker satisfaction for each utilization (0.1).

We index cost, duration, transport vehicle utilization, layout crew utilization, packing crew utilization, and signal crew utilization as one through six, respectively. Cost is selected as the standard measure and we assess $\delta_1^* = 0.00434$ by inverting (3), anchoring the cost at \$120 thousand, and then assessing an indifference zone value in the positive direction of \$1,000. From Proposition 2, the equivalent indifference zone parameters for all the other dimensions are:

$$\delta_2^* = 0.00869, \delta_3^* = \delta_4^* = \delta_5^* = \delta_6^* = 0.01737.$$

4.1 Sensitivity Analysis and Indifference Zone, Preference Zone Diagrams

Indifference-zone, preference-zone diagrams are constructed using *certainty equivalents (CEs).* For a single attribute utility function, the certainty equivalent is equal to the inverse of the utility function evaluated at the expected utility (Clemen 1991, page 372), i.e.,

$$\mathrm{E}[u_i(X_{ki})] = u_i(CE_{ki})$$

for $1 \leq i \leq n$, $1 \leq k \leq K$.

The curve dividing the indifference-zone from the preference-zone is constructed by setting

$$u_i(CE_{[K]i}) - u_i(CE_{[K-1]i}) = \delta_i^*$$

and solving for $CE_{[K]i}$. For the utility function in (1), the resultant expression is

$$CE_{[K]i} = CE_{[K-1]i} + RT\ln\left\{-\left(\frac{\delta_i^*}{B}\right)e^{-\left(\frac{CE_{[K-1]}}{RT}\right)} + 1\right\}.$$

For the cost measure with $\delta_1^* = 0.00434$, the indifference zone, preference zone diagram is given in Figure 2 The indifference (preference) zone lies above (below) the curve.

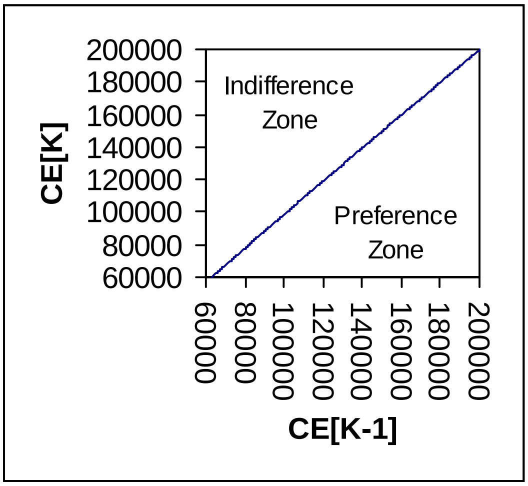

Figure 2: Indifference Zone, Preference Zone Diagram for Cost

To illustrate how indifference zone, preference zone diagrams can be used in sensitivity analysis, we will vary the weight on the cost performance measure (i.e., w_1) while holding the ratios amongst the other weights constant. This is a common approach to one-dimensional sensitivity analysis used in MAU theory (Butler et al. 1997). Using Proposition 2, Figures 3 and 4 contain indifference, preference zone diagrams for job duration and transport vehicle utilization as w_1 is varied from *0.1* to *0.9* in increments of *0.4*. Note that indifference zone, preference zone diagrams for all other utilizations are the same as Figure 4 since Equation (4) is used as the single attribute utility functions in all cases. In both figures the indifference zones always lie above (or to left of) the curves and the preference zones lie below (or to the right of) the curves.

For both performance measures, the curves in Figures 3 and 4 shift to the right and assume a greater slope and curvature as w_1 increases. The explanation for this rests on the fact that as the w_1 increases, the weight on the other two measures decreases, i.e., job duration and transport vehicle utilization are less important relative to cost. Therefore, larger changes in job duration and transport vehicle utilization are required to be equivalent to a given change in cost. For example, with the indifference zone parameter $\delta_1^* = 0.00434$ for cost, with $w_1 = 0.9$, $\delta_2^* = 0.11725$, $\delta_3^* = \delta_4^* = \delta_5^* = \delta_6^* = 0.2345$ (compare these values with values given above when $w_1 = 0.4$). The larger δ-values yield indifference zone curves with greater slope and curvature.

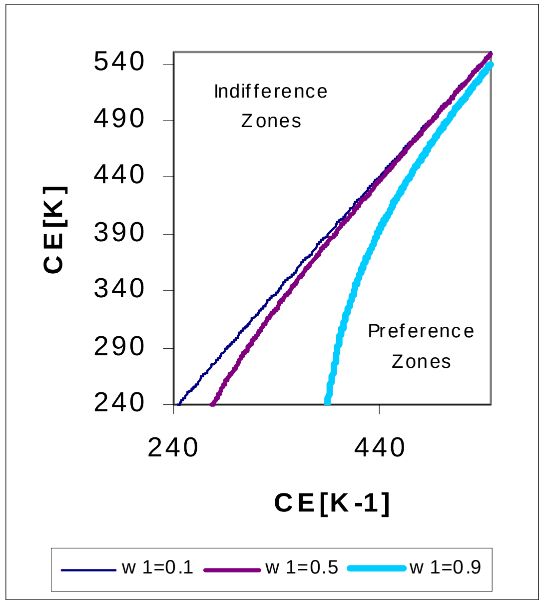

Figure 3: Indifference Zone, Preference Zone Diagrams for Job Duration

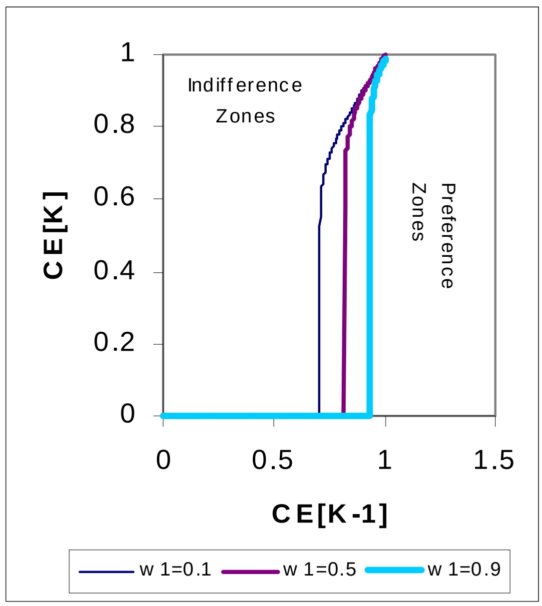

Figure 4: Indifference Zone, Preference Zone Diagrams for Transport Vehicle Utilization

In this manner, the indifference zone, preference zone diagrams can be used for sensitivity analysis and to check the consistency of the decisions maker's preferences reflected in the weights.

4.2 Impact of Sensitivity Analysis on R&S Results

Once again, we will perform sensitivity analysis on the weights by changing one weight at-a-time while holding the ratios amongst the other weights constant. Table 1 contains sensitivity analysis results for w_1. The range of variation from *0.25 to 0.9* reflects the importance of the cost parameter. It is doubtful that a decision maker would ever want to put less that twenty-five percent of the weight on cost and conceivable that (s)he would put almost all the weight on cost (i.e., ninety percent). The column label "Stat." contains a row for the mean utility for the second stage R&S results ("Mean") and a row for the total number of simulation runs ("#Sim.) for each value of w_1 considered.

The results in Table 1 illustrate that the robustness of Configuration 2 as the best configuration (highest mean utility) over all values of w_1 considered. Likewise, Configuration 4 remains a robust performer as the second best configuration.

With the fewest resources, Configuration 1 has longest duration. As a result, Configuration 1 experiences rapid improvement as w_1 increases because weight shifts from the project duration to the project cost. Additionally, the number of simulation runs drop for the same reason because project duration has a high level of variability. (It turns out that the variability in utility of job duration is at least one order of magnitude greater than the variability in the utilities on the other measures for Configuration 1.)

Configuration 3 has one more source crew than Configuration 1 so, on average, the former has a shorter duration than the latter. However, 1100 units of equipment rather than one source crew is the bottleneck so the additional source crew in Configuration 3 makes this configuration more expensive than Configuration 1. Both of these factors explain why Configuration 3 performs better than Configuration 1 for lower values of w_1 but the latter improves more quickly than the former as w_1 increases.

The number of simulation runs required for Configuration 3 decreases because at w_1 = 0.3 the weights on costs and project duration are roughly the same (w_2 = 0.233 and all the other weights equal 0.117). Additionally, the utilities of the two performance measures have roughly the same variance for this configuration (which is at least three orders of magnitude greater than the variances associated with the utilities of the other performance measures for this configuration). Therefore, the variance of $u_1(X'_{31})$ is determined by the sum of the variances of the cost and project duration utilities (see the formula in Step 2 of the procedure in Section 3) plus a covariance term which is positive since these two measures are highly positively correlated. As w_1 increases, the variance of $u_1(X'_{31})$ approaches the variance of the utility of the cost (i.e., something less than half of the variance of $u_1(X'_{31})$ for w_1 = 0.3).

Table 1: R&S Results from Cost Weight Sensitivity Analysis

w_1	Stat.	Configuration			
		1	2	3	4
0.25	Mean	0.270	0.931	0.419	0.782
	#Sim.	820	100	400	130
0.3	Mean	0.402	0.944	0.473	0.818
	#Sim.	520	100	340	100
0.35	Mean	0.501	0.953	0.511	0.842
	#Sim.	350	100	300	100
0.4	Mean	0.576	0.960	0.540	0.861
	#Sim.	250	100	280	100
0.45	Mean	0.631	0.965	0.562	0.875
	#Sim.	180	100	260	100
0.5	Mean	0.676	0.969	0.579	0.886
	#Sim.	130	100	240	100
0.55	Mean	0.711	0.972	0.594	0.896
	#Sim.	100	100	230	100
0.6	Mean	0.742	0.975	0.607	0.904
	#Sim.	100	100	230	100
0.65	Mean	0.768	0.978	0.617	0.910
	#Sim.	100	100	220	100
0.7	Mean	0.791	0.980	0.626	0.916
	#Sim.	100	100	220	100
0.75	Mean	0.811	0.982	0.634	0.921
	#Sim.	100	100	210	100
0.8	Mean	0.828	0.983	0.640	0.925
	#Sim.	100	100	210	100
0.85	Mean	0.843	0.985	0.646	0.929
	#Sim.	100	100	210	100
0.9	Mean	0.856	0.986	0.652	0.932
	#Sim.	100	100	200	100

We chose 0.1 to 0.4 as the most realistic range of variation for w_2. Table 2 contains sensitivity analysis results. Once again, Configurations 2 and 4 remain the best and second best, respectively, over the entire range of variation. With an increased emphasis on job duration, Configuration 1 deteriorates rapidly and requires more observations for the reasons given above. Additionally, the same arguments used above explain the results for Configuration 3 in Table 2.

Table 2: R&S Results from Job Duration Weight Sensitivity Analysis

w_2	Stat.	Configuration			
		1	2	3	4
0.1	Mean	0.742	0.965	0.607	0.863
	#Sim.	100	100	230	100
0.15	Mean	0.664	0.962	0.575	0.862
	#Sim.	140	100	250	100
0.2	Mean	0.576	0.960	0.540	0.861
	#Sim.	250	100	280	100
0.25	Mean	0.471	0.956	0.500	0.859
	#Sim.	400	100	310	100
0.3	Mean	0.356	0.953	0.454	0.857
	#Sim.	620	100	360	100
0.35	Mean	0.224	0.949	0.399	0.855
	#Sim.	940	100	430	100
0.4	Mean	0.065	0.944	0.335	0.853
	#Sim.	1390	100	520	100

We performed one-at-a-time sensitivity analysis on w_i, $i = 3, 4, 5, 6$ over a range deemed realistic: 0.05 to 0.25. Only the sensitivity analysis on w_4 proved to be interesting because the layout crew becomes a bottleneck after the amount of receiving equipment is increased to 1300 units. All other crews have low utilization in all four scenarios. Hence sensitivity analysis on their weights had almost no impact on the results.

Table 3 contains results for sensitivity analysis on w_4. Changes in w_4 over the range *0.05* to *0.25* do not affect Configurations 1 and 3 since equipment constraints production in both of these configurations and not the layout crew. For Configurations 2 and 4, the layout crew becomes the constraining resource. As a result, both of these configurations begin to deteriorate as w_4 increases to 0.25. However, they still remain the top two configurations in all cases given in Table 3. Further increases in w_4 lead to further deterioration of these two scenarios. We elected not to include such results in Table 3 because, in practice, it is unlikely that a decision-maker would put more than twenty-five percent of the weight on a single crew utilization measure. The increased number of simulations for Configuration 4 when w_4 = 0.25 results from the relatively high the variance of the utility for layout crew utilization (the variance of the utility for the layout crew utilization is at least four times the variance of the other utilities for this configuration).

In combination, the results from Tables 1, 2 and 3 indicate that Configuration 2 represents a robust best configuration. However, we must reemphasize that the sensitivity analysis results are limited by the fact that we have only considered one-at-a-time changes in the weights.

Table 3: R&S Results from Layout Crew Utilization Weight Sensitivity Analysis

w_4	Stat.	Configuration			
		1	2	3	4
0.05	Mean	0.576	0.969	0.540	0.898
	#Sim.	250	100	280	100
0.1	Mean	0.576	0.960	0.540	0.861
	#Sim.	250	100	280	100
0.15	Mean	0.576	0.949	0.540	0.819
	#Sim.	250	100	280	100
0.2	Mean	0.576	0.936	0.540	0.772
	#Sim.	250	100	280	120
0.25	Mean	0.576	0.922	0.540	0.715
	#Sim.	250	100	280	210

5 CONCLUSION

In this paper, we have conducted a one-at-a-time sensitivity analysis on the weights of a combined MAU theory R&S procedure. Future research includes sensitivity analysis by making simultaneous changes in the weights using approaches similar to those found in Butler et al. (1997). Additionally, we plan to conduct sensitivity analysis on the single attribute utility functions.

ACKNOWLEDGEMENTS

The authors would like to thank Peter Highnam of Schlumberger Austin Product Center for reviewing and supporting this work.

REFERENCES

Butler, J., J. Jia, and J. Dyer. 1997. Simulation Techniques for the Sensitivity Analysis of Multi-Criteria Decision Models. *European Journal of Operational Research* 103: 531-546.

Bechhofer, R. E., T. J. Santner, D. M. Goldsman. 1995. *Design and Analysis of Experiments for Statistical Selection, Screening, and Multiple Comparisons*. New York: John Wiley & Sons, Inc.

Clemen, R. T. 1991. *Making Hard Decisions.* Boston: PWS Kent Publishing.

Goldsman, D. M., B. L. Nelson, B. W. Schmeiser. 1991. Methods for Selecting the Best System. In *Proceedings of the 1991 Winter Simulation Conference,* ed. B. L. Nelson, W. D. Kelton, and G. M. Clark, 177-186. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

Keeney, R. L. and H. Raiffa. 1976. *Decisions with Multiple Objectives.* New York: Wiley.

Morrice, D. J., J. C. Butler, and P. W. Mullarkey. 1999. A Multiple Attribute Utility Approach to Ranking and Selection for Systems With Multiple Performance

Measures. Working Paper. The University of Texas at Austin.

Morrice, D. J., J. C. Butler, and P. W. Mullarkey. 1998. An Approach to Ranking and Selection for Multiple Performance Measures. In *Proceedings of the 1998 Winter Simulation Conference***.** D. J. Medieros, E. F. Watson, J. S. Carson, and M. S. Manivannan, eds. The Institute of Electrical and Electronic Engineers, Piscataway, New Jersey, 719-725.

Morrice, D. J., P. W. Mullarkey, A. S. Kenyon, H. Schwetman, J. Zhou. 1997. Simulation of a Signal Quality Survey. In *Proceedings of the 1997 Winter Simulation Conference,* ed. S. Andradottir, K. J. Healey, D. H. Withers, and B. L. Nelson, 1265-1272. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

Rinott, Y. 1978. On Two-Stage Procedures and Related Probability-Inequalities. *Communications of Statistics: Theory and Methods* A8:799-811.

AUTHOR BIOGRAPHIES

DOUGLAS J. MORRICE is an Associate Professor in the MSIS Department at The University of Texas at Austin. He has a BA, Honours in Operations Research from Carleton University in Ottawa, Canada. He holds an M.S. and a Ph.D. in Operations Research and Industrial Engineering from Cornell University. His research interests include operations simulation design, modeling, and analysis. Dr. Morrice is a member INFORMS. He served as the Secretary for the INFORMS College on Simulation (1994-1996) and was Co-Editor of the *Proceedings of the 1996 Winter Simulation Conference*.

JOHN BUTLER is a Ph.D. candidate in the College of Business Administration at the University of Texas at Austin. His research interests include decision theory, decision support systems, and mathematical modeling. Currently he is involved in a research project designed to aide in the selection of a technology for the disposition of US surplus weapons grade plutonium.

PETER W. MULLARKEY is a Research Scientist in the End-to-End Simulation Department at Schlumberger Austin Product Center - Research. He received his undergraduate degree in Civil Engineering at Worcester Polytechnic Institute in Worcester, Massachusetts. He holds a Ph.D. in Civil Engineering from Carnegie-Mellon University. Dr. Mullarkey has worked with Schlumberger doing research and development of knowledge-based decision support systems since 1985. Additionally, he worked on a two year project at Tandem Computers developing a state model-based diagnosis and maintenance system product. In April 1999, Dr. Mullarkey left Schlumberger to join Maxager Technology (www.maxager.com), a startup based in San Rafael, CA, whose focus is to use advanced software to help manufacturers make tactical and strategic decisions based on the metric of profit-per-minute.

SRINAGESH GAVIRNENI is a research scientist in the Austin Product Center of Schlumberger. He has an undergraduate degree in Mechanical Engineering from Indian Institute of Technology, Madras, India. He also has a Masters in Industrial Engineering from Iowa State University and a Ph.D. in Operations Management from Carnegie Mellon University. His research interests are in the areas of Supply Chain Management, Inventory Control, Scheduling, and Resource Management. His papers have appeared in Management Science, Manufacturing & Service Operations Management, European Journal of Operational Research, and IEEE Transactions on Reliability.