PERFORMANCE INDICES TO AID NUCLEAR MATERIAL SAFEGUARD MANAGEMENT DECISIONS

John W. Lathrop

August 1981
PP-81-12

INTERNATIONAL INSTITUTE FOR APPLIED SYSTEMS ANALYSIS
2361 Laxenburg, Austria

ACKNOWLEDGEMENTS

Many of the thoughts presented in this paper originated in discussions concerning related work while the author was at Lawrence Livermore Laboratory. That research was done for the Nuclear Regulatory Commission, with Woodward-Clyde Consultants and Decisions and Designs, Incorporated, as subcontractors. An acknowledgement is due for the stimulating conversations and helpful comments of several people from those organizations, in particular Rex Brown (now at Decision Science Consortium), Ralph Keeney (Woodward-Clyde Consultants), William Murphey (now at the Arms Control and Disarmament Agency), and Leo Spogen and Stein Weissenberger (both at Lawrence Livermore Laboratory). While these people helped form many of the ideas presented here, the author accepts full responsibility for the contents of this paper.

CONTENTS

PERFORMANCE INDICES TO AID NUCLEAR MATERIAL SAFEGUARD MANAGEMENT DECISIONS

John W. Lathrop

I. INTRODUCTION

The management of a national or international nuclear material safeguard system involves large numbers of very complex strategic and tactical decisions at several organizational levels. This report outlines a decision analytic approach to developing decision aids to assist the various decision makers in the material accounting part of that process. The intent of the report is to lay out a very general, simple conceptual framework for such decision aids. While the framework combines the most basic ideas of decision analysis, game theory, and signal detection theory, the nature of the approach has been kept at an easily understandable level, in order to enhance its chances of implementation in the real world of nuclear safeguard management. As a result, this framework does not begin to approach the thoroughness and sophistication of recent statistical treatments of the problem, such as the work of Avenhaus (1977). Instead, the approach presented here emphasizes a simple conceptual structure built around the basic idea set out in Bennet, Murphey, and Sherr (1975): safeguard management should be guided by some measure of the societal risk or social utility that results from the operation of the safeguard system.

This document is a status report, representing a very small number of person-days of effort. It describes a very "quick and dirty" first cut at the problem, to give the people who would be involved in the implementation of these decision aids an opportunity to criticize and comment at a very early stage.

As is typical of a decision analytic approach, the suggested decision aids are built around a set of performance indices. Section II presents the roles the set of indices would play in the safeguard management process, while Section III lays out the nature of decision problems involved. Then in Section IV a particular set of indices is suggested and described. Section V discusses some of the most significant problems to be confronted in implementing any set of decision aids. The paper concludes with a suggested agenda for developing, testing, and implementing the decision aids.

II. ROLES OF PERFORMANCE INDICES

In the most general way, a system of performance indices developed within a decision analysis framework can serve to bring a certain rationality, or at least self-consistency, to a system that might otherwise exhibit inconsistencies often found in ad hoc, bureaucratic and political systems. While this general role may in fact result in the most important contribution of a set of performance indices, this section will briefly describe more specific roles of the indices. Performance indices can aid safeguard management decisions at each of three levels:

- low-level allocation of inspection resources;
- higher-level evaluation of safeguard approaches and system-wide efficiency testing;
- assisting in highest-level determination of adequacy of the overall safeguard system.

Because each of these decision levels can involve a different type of index, they are treated separately in the following subsections.

A. Allocation of Inspection Resources

The allocation of inspection resources can best be guided by an importance index over diversion possibilities (*DP*'s) that indicates the relative importance of safeguarding each *DP*. Such an index could be used:

- to aid decisions involving the distribution of frequency and type of inspections;

- to develop instructions to inspectors;
- to set inspection goals.

B. Safeguard Approach Evaluation

The evaluation of safeguard approaches calls for a performance index that aggregates the effectiveness of the safeguards over all known *DP*'s in each nuclear facility. Such an index could be used to evaluate:

- generic safeguard approaches,
- actual, facility-specific safeguard approaches,
- particular inspector actions.

Once a performance index is developed, it can be used to seek efficient allocation of resources by guiding re-allocations of safeguard effort to maximize the index for a given budget.

C. Aiding the Determination of Safeguard Adequacy

The evaluation of the entire national or international safeguard system requires a performance index that aggregates the effectiveness of safeguards over all safeguarded facilities. Such an index could be used in a determination of whether or not the overall safeguard budget is large enough for the system to perform adequately. It would be possible to develop an index that determines adequacy directly. That index would incorporate the crucial cost-vs-safety tradeoff within the inner workings of its calculations. However, it would probably be much more defensible to develop an index that measures overall system performance in terms of resulting safety aside from its cost, then leave the cost-vs-safety debate to the political process. In order to be the most helpful for such a debate, the index would express the level of safety resulting from the system in terms meaningful to the participants in that political process, such as an overall probability of detection given diversion of one significant quantity (SQ).

III. THE NATURE OF THE PERFORMANCE EVALUATION PROBLEM

The development of performance indices requires that a series of very difficult questions be asked concerning just how performance should be evaluated, and just what is it that constitutes good performance. This section presents some of those questions and considerations.

A. Ultimate Value Dimensions

The performance of a safeguard system is multidimensional. In addition, there are several levels at which the safeguard system can be described multidimensionally. At the lowest level, which we can call level N, the descriptive dimensions are the observables: frequency and types of inspections, numbers of anomalous situation indications (ASI's), camera performance, etc. While it may be relatively easy to describe a safeguard system on these directly observable dimensions, such a description is far removed from what the managers of the safeguard system and its participant organizations actually care about.

At the highest level, which we can name Level I, performance description is relatively straightforward: whether or not a nuclear weapon made with diverted material is used for warfare, political pressure, or economic blackmail. But that level is too high for material accounting safeguard performance indices, since it involves a much larger system of political checks and balances than the material accounting safeguard system alone, whose function is simply to "ring the bell." One step down from this level, Level II, we can identify more narrowly defined, more easily measurable description dimensions. One especially simple example set of dimensions at Level II can be listed if one assumes that no more than one diversion and no more than one alarm could occur per year:

- annual probability of correct detection before a weapon is built (early correct alarm, ECA),
- annual probability of correct detection, but not until after a weapon is built (late correct alarm, LCA),
- annual probability that a diversion of more than one SQ occurs without detection (false rejection, FR),
- annual probability that a diversion is reported that did not in fact occur (false alarm, FA).

Given the above assumptions, the complement of the sum of these probabilities is the annual probability the system remains correctly silent (correct rejection, CR). A more complete set of dimensions would break each probability down by amount diverted, and would allow for multiple diversions and alarms per year. All of the above probabilities are annual, and are absolute, i.e., not conditioned on a diversion taking place. This absoluteness is desirable in that it allows the performance description to include the safeguard system's power to deter a diversion attempt. However, such absolute probabilities cannot be calculated without involving debatable assumptions that may degrade the credibility of the performance descriptions. It may be desirable, on balance, to step down one more step in performance description, to Level III, and change the upper three probabilities to be conditional on a diversion. The descriptive dimensions at Levels II and III are very close to what the safeguard managers (and the society they are protecting) care about concerning safeguards, and so can be referred to as value dimensions.

This discussion has the effect of outlining the requirements for the process models that must support a system of performance indices: Those models must take as inputs the observables of Level N and deliver as outputs a set of descriptors at Level III or Level II.

B. Value Tradeoffs and Attitude Toward Uncertainty

The previous subsection described sets of four value dimensions that can be used to describe the performance of a safeguard system. Certainly the four descriptors should be presented as a set to the appropriate decision makers reviewing the performance of the system. However, any decisions involving the comparison of two safeguard approaches or suggested changes in the system must involve the aggregation of the four (or whatever number) of descriptors into a single index of overall performance, either implicitly or explicitly. Without such a single index, one cannot say "system A is better than system B," or even "system A is efficient," strictly speaking.

The step from several descriptors to one index is a step from description to evaluation. That step cannot be taken without making some value tradeoffs between the value dimensions, and incorporating some attitude toward uncertainty. For example, suppose a reallocation of inspection resources is being considered that would increase inspection frequency on some activities and drop some other inspection activities. At Level III, the decision on whether or not to make that reallocation can be represented as a choice between sets of probabilities such as the following (all conditional on "div," a diversion taking place):

p(ECA|div)	p(LCA|div)	p(FR|div)	p(ECA|div)	p(LCA|div)	p(FR|div)
{ .1	.6	.3 }	{ .4	.25	.35 }

The choice as to whether the left or right set is preferable requires value tradeoffs and some particular attitude toward uncertainty.

Process models would allow the mapping of value tradeoffs and attitude toward uncertainty at Level III down into lower levels, where they may have more meaning for tactical decisions. Two examples of tactical decisions requiring value tradeoffs:

1) You have two unsafeguarded diversion possibilities (*DP*'s), one involving 16 kg Pu, the other 8 kg Pu. For the amount remaining in your budget, you can get an 80% detection probability on the 16 kg *DP*, or a 99% detection possibility on the 8 kg *DP*. Which do you choose?

2) Suppose DP_1 involves 1 SQ, and DP_2 involves 9 SQ, yet DP_1 is 9 times more likely to be attempted than DP_2. Both have a detection probability of 30%. For the amount remaining in your budget, you can increase the detection probability in one or the other to 40%. Which do you choose?

The value tradeoffs and uncertainty attitude involved in the above dilemmas do not have an agreed-upon objective basis. They represent social value tradeoffs that must be made, implicitly or explicitly, in the course of managing a safeguard system. A set of performance indices would incorporate the less crucial values in a structured, systematic way, using tradeoffs and uncertainty attitudes elicited from the appropriate decision makers by means of accepted decision analytic techniques. The most crucial value tradeoff, cost-vs-safety, can be left out of the index system, to be decided by the appropriate political debate. The index system can, however, help to structure that debate and couch it in the performance evaluation terms most meaningful for the participants.

C. Structure Between Evaluation and Safeguard Management Decisions

There is a basic structure linking the value tradeoffs discussed in the previous section with the safeguard management decisions to be aided by the performance indices. That structure is most clearly presented in the form of a decision tree, as shown in Figure 1. The square nodes in the figure represent decision points, manned from left to right by the inspector, his supervisors, then the highest-level safeguard manager. The circular nodes represent random events. For example, if the inspector rejects a signal as not significant, it may either be a false rejection of a real diversion, or a correct rejection, with a positive probability for each. For clarity, the CA in Figure 1 represents a correct alarm, whether it is early or late.

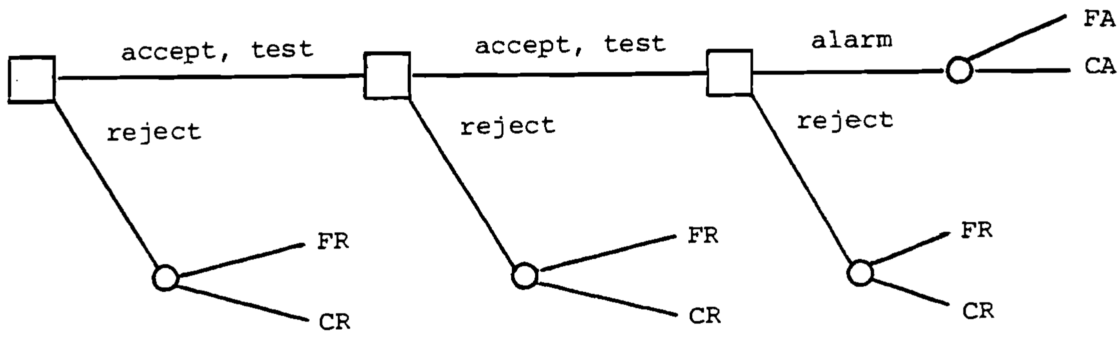

Figure 1.

The decision tree in Figure 1 displays the ultimate outputs of a safeguard system (tightly linked to a Level III performance description) and how they are linked to signal detection decisions. Safeguard management decisions affect the discriminability of the signal detection systems used by each decision maker (d' in signal detection theory (SDT) notation), and can affect the criteria used to accept or reject a signal (β in SDT notation, a quantity different from the β of safeguard terminology. See Egan 1975). Section IV presents a system of indices which evaluate probability distributions over the safeguard outputs indicated in Figure 1 in ways that can be used as very direct guidance for safeguard management decisions.

D. Requirements of a System of Performance Indices

A system of performance indices must fulfill the roles outlined in Section II, incorporating the value dimensions and tradeoffs mentioned in subsections A and B above in the structure presented in subsection C above. But beyond all that, there are requirements that the index system must meet in order to be feasible and useful. These are listed here.

1. The index system must treat the diverter as a gamesman.

A system of performance indices must treat the diverter as one who will react to the safeguard system to maximize his own gain. Any index system that treats the diverter as if he is blind to safeguard measures will have serious problems with defensibility.

2. The index system must be understandable.

In order for the performance index system to be useful, it must be accepted and easily understood by all those who use it. That means the system must be relatively straightforward, with a clear and logical structure, even at the cost of some decision-theoretic rigor.

3. *The index system must be compatible with a feasible and realistic inspection and data gathering system.*

The index system must be able to use the inspection system as it currently operates. It must also be able to identify feasible and politically realistic improvements in that system.

IV. A SUGGESTED SYSTEM OF PERFORMANCE INDICES

The previous sections have presented guidelines and requirements that should be fulfilled by any system of performance indices. In this section, one suggested system is put forward as an illustration of what an index system might look like.

A. Overall Structure

In keeping with the three levels of roles presented in Section II, the suggested index system has three levels:

1. Diversion Possibility Level: Games Against the Diverter

At this level, a utility function for the diverter, U_D, is used as an importance index, attaching a measure of diverter attractiveness to each DP_i given safeguard approach SA_j : U_D $(DP_i \mid SA_j)$. This function takes into account the probability of detection and diversion success, as well as value to the diverter of the various possible outcomes and the technical difficulty, complexity, and cost to the diverter. In allocating detection resources among DP_i's, one simple principle is to select the SA_j that makes the $U_D(DP_i \mid SA_j)$ values as equal as possible across DP_i's. The reasoning here is that the diverter will strike at the "weakest link" in the safeguard defense, which is the DP_i most attractive to him (i.e., with the highest $U_D(DP_i \mid SA_j)$). By this reasoning, detection resources are best spent in decreasing the attractiveness of the DP_i with the highest $U_D(DP_i \mid SA_j)$..

While this level of index is intended to guide allocation of detection resources over detection possibilities, that does not provide direct guidance for allocation of inspection resources. A paper by Ulvila (1980) describes the relationship between DP's and inspection activities, and in doing so, casts some light on the problems of translating resource allocations over DP's into allocations of inspection activities.

2. Safeguard Approach Level: Seeking Efficiency

At this level, a utility function for society, U_S, is developed as an index of how well society is served on the Level III value dimensions by the safeguard approach, SA_j. The U_S (SA_j) can be used to seek efficiency by reallocating inspection resources and adjusting alarm threshold points until U_S (SA_j) is maximized for a given budget level.

3. Safeguard System Adequacy Level: Aiding the Debate

The U_S index can be applied at a system-wide level to measure the maximum achievable system performance for each of several budget levels. These performance-cost pairs can then be presented to the political decision making process for a decision on which pair to implement (i.e., which budget level to set). Rather than present the performance of the system in the rather abstract units of U_S, a more meaningful numeraire could be used, such as the overall probability of detection given diversion of one SQ that corresponds to that value of U_S. This process very carefully leaves the cost-vs-safety tradeoff in the domain of political process.

The relationship between efficiency and adequacy determinations is illustrated in Figure 2. In the figure, a carefully defined probability of detection is used as a meaningful numeraire for U_S. The dots represent

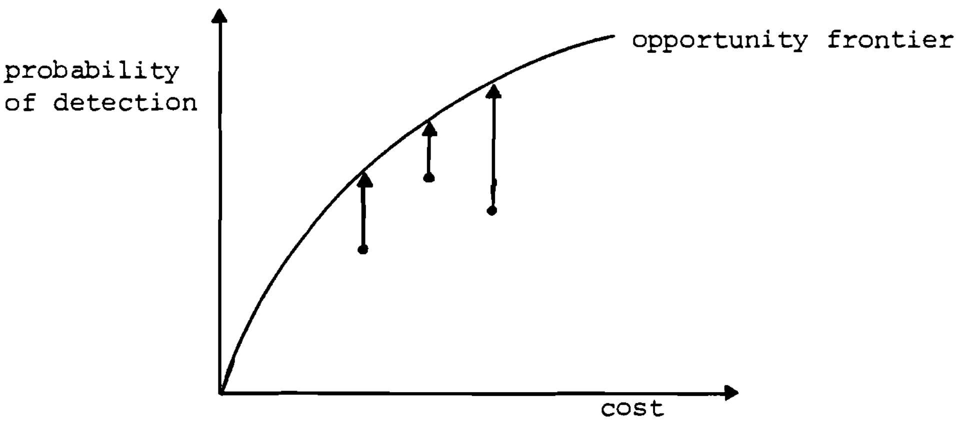

Figure 2.

originally proposed safeguard systems, at three different budget levels. The arrows indicate the adjustment of each of those systems to maximize U_S subject to the budget constraint. The opportunity frontier represents all possible efficient performance-cost pairs. The entire frontier would never actually be plotted. The political process would debate only over which of the three efficient systems to select. The outcome of each year's budget debate could be used to guide the selection of the subsequent year's three trial systems.

B. Utility for the Diverter, U_D

The diverter cannot know before the attempt what its outcome will be. His choice among *DP*'s is a case of decision making under uncertainty, where risk aversion may be important. An appropriate way to represent his preference, then, is with a von Neumann-Morgenstern utility function, u_d, over the possible outcomes. His utility for a diversion possibility, $U_D(DP_i \mid SA_j)$, is then the expectation over his utilities of the possible outcomes. The outcome utility function for the diverter, u_d, can be considered a function of the number of SQ's diverted, when, if ever, he is detected, and the technical difficulty, cost and complexity of the *DP*. The utility function should not represent the preferences of a normative diverter, but should be as descriptive as possible of actual diverter choice behavior. With this in mind, the utility function should be additive, since there is ample experimental evidence of the descriptive validity of such a preference model (Dawes and Corrigan 1974). Beyond this question of structure, wherever possible the u_d function should be elicited by wholistic judgments, again to maintain descriptive validity.

In order to keep the u_d function as simple as possible, while still capturing the essential aspects of the situation it is to describe, its argument space could be numbers of SQ diverted (denoted X) crossed with three detection times (denoted {before, after, never} for before the diverter has built a weapon, after that, and never) crossed with perhaps four levels of technical difficulty-complexity-cost, described by paragraphs or simply by equivalent monetary cost (denoted C, with C_i the cost of DP_i).

In sum, then, the diverter's utility for a diversion possibility given a safeguard approach, $U_D(DP_i \mid SA_j)$, is his expected utility over the uncertainty the diverter faces: whether or not he will succeed in getting x SQ's and when if ever he will be detected. Equation (1) displays that $U_D(DP_i \mid SA_j)$ for the simple case when amount diverted can only be x or zero:

$$\begin{aligned} U_D(DP_i \mid SA_j) = \quad & p\ (x, never \mid DP_i, SA_j) \cdot u_d\ (x, never, C_i) \\ & + p\ (x, after \mid DP_i, SA_j) \cdot u_d\ (x, after, C_i) \\ & + p\ (x, before \mid DP_i, SA_j) \cdot u_d\ (x, before, C_i) \qquad (1) \\ & + p\ (o, never \mid DP_i, SA_j) \cdot u_d\ (o, never, C_i) \\ & + p\ (o, after \mid DP_i, SA_j) \cdot u_d\ (o, after, C_i) \\ & + p\ (o, before \mid DP_i, SA_j) \cdot u_d\ (o, before, C_i). \end{aligned}$$

C. Utility for Society, U_S

At the time a safeguard approach must be selected, there is substantial uncertainty concerning the magnitude of the diversion threat and the effectiveness of that approach in meeting the threat. As was the case with the diverter's choice among DP's, then, a choice among safeguard approaches is a case of decision making under uncertainty, where risk aversion may be important, so an appropriate way to represent society preferences over approaches is with a von Neumann-Morgenstern utility function, u_s, over the possible outcomes. The utility for society of a safeguard approach, $U_S(SA_j)$, is then the expectation over the societal utilities of the possible outcomes. Up until this point there have been marked similarities between U_D and U_S, and between u_d and u_s. However, the outcome utility function for society, u_s, is basically different from u_d in that it should be prescriptive, not descriptive. To take the most simple case, the u_s argument space could be amount diverted, X, crossed with a five-value performance variable, which we will call R (for "response"), taking values {CR, ECA, LCA, FR, FA} corresponding to those responses discussed in terms of the Level II and III description dimensions presented in Section III. In keeping with the Level III dimensions, the utility for society for a safeguard approach, SA_j, or even for the entire safeguard system can be an expectation conditional on a diversion ("div") taking place,

$$U_S(SA_j) = \sum_{k,l} p\ (x_k, r_l \mid SA_j,\ div) \cdot u_s\ (x_k, r_l), \tag{2}$$

where

$$p\ (x_k, r_l \mid SA_j,\ div) = \sum_i p\,(x_k, r_l \mid SA_j,\ DP_i) \cdot p\ (DP_i \mid SA_j,\ div), \tag{3}$$

and where

$$p\ (DP_i \mid SA_j,\ div) = \frac{U_D\ (DP_i \mid SA_j)}{\sum_i U_D\ (DP_i \mid SA_j)}\ . \tag{4}$$

Equation (4) is borrowed from the experimental psychology literature on stochastic choice, and is commonly referred to as the Luce choice model (1959). While more elaborate choice models could be employed, in view of the great uncertainties involved in the formulation of U_D, a large choice modeling effort would probably not be justified.

The astute reader will note that the idea of equating the $U_D(DP_i \mid SA_j)$ over DP_i, presented in Section IV.A.1 as a way to game against the diverter, is a technically superfluous feature of the proposed optimization process. This is because optimizing $U_S(SA_j)$ will in fact allocate resources across DP_i in a desirable way, accounting for differences in $p\ (x_k, r_l \mid SA_j, DP_i)$ and $u_s\ (x_k, r_l)$, as well as differences in $U_D\ (DP_i \mid SA_j)$, and allowing for some noise in the diverter's choice behavior. The equating of $U_D\ (DP_i \mid SA_j)$'s is retained as an optional, separate guide to low-level resource allocation because its simplicity and logic make it a more implementable policy guide than the more obscure Equations (2) through (4), especially in a decentralized safeguard system.

An additional use of u_s (technically contained in the above equations but hardly apparent by inspection) is in the setting of threshold points for the accept-reject decisions represented in the decision tree of Figure 1. To make a clear example, consider the last decision node on the right in that tree, assume only one possible diverted amount x, and consider early and late correct alarms as equivalent (denoted CA). It is easily shown from signal detection theory (Egan 1975), or more directly by simply equating expected utilities, that the evidence should be accepted as indicating an alarm should be sounded when the probability of diversion given the evidence, $p\ (x \mid evid)$ is such that

$$\frac{p\ (x \mid evid)}{1-p\ (x \mid evid)} \geq \frac{u_s\ (o, CR) - u_s\ (o, FA)}{u_s\ (x, CA) - u_s\ (x, FR)} \quad . \tag{5}$$

Naturally, more complicated expressions are required for more than one possible amount diverted, for when the detection time (early vs late) matters, and for the intermediate detection decision stages. However, the basic idea stands that the optimal accept-reject threshold point is a function of the social utilites of the possible outcomes.

The last paragraph serves to illustrate the pervasive usefulness of this system of performance indices. It can be used to decide whether to use more funds to increase discriminability and lower the false alarm rate, or to increase the detection rate, or to use the funds to reduce detection time. It can be used to optimize the staging of the detection system illustrated in Figure 1. It can be used under politically set constraints, such as a maximum false alarm rate, to find the optimal settings of the unconstrained parameters. Then it can be used to measure the social cost of that politically set constraint!

D. Information Requirements

Perhaps the most important thing to remember about this performance index system is that it does not by itself generate any guidance for safeguard management. All it does is provide a logical, self-consistent framework for moving from collecting the required basic information to making the required decisions. There are two basic types of information called for: value information and technical information.

1. Value Information

The u_d and u_s functions must be elicited from the appropriate individuals. There is a well-established methodology for doing this elicitation (Keeney and Raiffa 1976), but a problem does remain in identifying and reaching the appropriate individuals. One does not need to elicit u_d from an actual diverter, but simply from an expert who could represent the preferences of a diverter. While the elicitation of value information may seem somewhat speculative in nature, it in fact is a strong feature of the suggested approach. Any approach that does not include such value elicitation must then assume some particular form for u_s and u_d, either explicitly or implicitly, without the benefit a structured value elicitation.

2. Technical Information

For each DP, a probability distribution must be estimated for the argument domain for u_s, X crossed with R. A methodology for eliciting such probability distributions from experts is described in Spetzler and Stael von Holstein (1975). In addition to this probabilistic information, a "cost" to the diverter, including technical difficulty and complexity, must be estimated for each DP. That cost can be expressed in monetary units or in a set of four or five paragraphs, each describing a particular level of difficulty in a general enough way to apply to all DP's.

V. PROBLEMS WITH PERFORMANCE INDICES

A. Dynamics of the Safeguard Process

An analysis of safeguard management as a static system may miss crucial aspects of the long-term safeguard process. As one example, consider the social preference (u_s) ranking of a false alarm versus a false rejection. A value elicitation in a static analysis may reveal a preference for the false alarm. Yet when the dynamics of the situation are taken into account, including the possible dismantling of the safeguard system in response to too many false alarms, it may be that in particular cases the false rejection is preferable. Clearly, this matter requires much more thought.

B. Differences in Values and Goals

Perhaps costs to the operator should be included in the analyses. However, it is not clear exactly how that cost would enter into U_S, and how much weight it should be given, relative to other arguments. In fact this particular dimension highlights a broader problem: what to do in cases where there is a clear difference between the values prescriptively ascribed to the society being protected, and the values of the safeguard system management agency. It could very well be the case, for example, that the institutional structure and incentives of the safeguard system agency causes its management to be more sensitive to operator cost than are members of the protected public. Of course, this problem is not one caused by the analysis. The problem has its roots in the institutional structure involved. However, the system of performance indices makes the values and goals of the system more explicit, and in doing so may draw attention to the problem.

C. Different Diverter Types

If an index system is to be responsive to the real world, it must acknowledge that there are different diverters out there. A system optimized against one diverter type is not optimized against the spectrum of possible diverters. The normative answer is to optimize the system against a probability distribution over diverter types. It would, of course, be quite difficult to obtain the necessary estimates of that distribution. One promising approach to alleviate this problem may lie in the development of a concept of resilience.

D. Information Loss in Aggregation

Any set of indices, by its very aggregative nature, must lose information. The point is to keep only that part of the information necessary to make management decisions. But particular features may be missed that are valuable in finding solutions to problems. One example was suggested by Frank Houk of the Arms Control and Disarmament Agency: suppose system performance plunges due to a flurry of reliability problems with cameras. The cause of the plunge may be lost in the indices, which only relay the plunge itself.

VI. FUTURE DIRECTIONS

The system of performance indices presented here exists in outline form only. There is much to do in order to implement such a system. The mathematics must be expanded to handle more general cases than those used as examples in this text. An actual large-scale safeguard system should be examined in some detail to investigate problems of compatibility between the indices and the system. Value elicitations should be performed, and the indices re-worked to handle surprises discovered in those elicitations. Then example and demonstration applications should be carried out, using utility functions fitted to the value elicitations. All of these steps should be carried out in close cooperation with the ultimate users of the indices in order to enhance the usefulness and acceptance of the indices. That cooperation can be continued past the methodological demonstrations to form the basis of the ultimate implementation of the indices.

The set of performance indices proposed here offers substantial promise as a system of decision aids for nuclear material safeguard management. Perhaps the strongest point of the approach is its simple conceptual structure, which promotes a ready understanding of the indices by its users and makes possible a realistic, multi-staged implementation strategy.

REFERENCES

Avenhaus, R. 1977. *Material Accountability: Theory, verification, and applications.* New York: Wiley.

Bennett, C.A., W.M. Murphey, and T.S. Sherr. 1975. "Societal Risk Approach to Safeguards Design and Evaluation," *ERDA* -7, June.

Dawes, R.M., and B. Corrigan. 1974. "Linear models in decision making," *Psychological Bulletin, 81.*

Egan, J.P. 1975. *Signal Detection Theory and ROC Analysis.* New York: Academic Press.

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

Luce, R.D. 1959. *Individual Choice Behavior.* New York: Wiley.

Luce, R.D., and H. Raiffa. 1957. *Games and Decisions.* New York: Wiley.

Spetzler, C.S., and C.-A.S. Stael von Holstein. 1975. "Probability encoding in decision analysis," *Management Science*, 22.

Ulvila, J.W. 1980. *A method for determining the best use of limited time available for the inspection of a nuclear facility.* (Technical Report 80-13). Falls Church, Virginia: Decision Science Consortium, Inc.