NOT FOR QUOTATION
WITHOUT PERMISSION
OF THE AUTHOR

DECISION ANALYSIS FOR THE EVALUATION OF RISK IN NUCLEAR WASTE MANAGEMENT

John W. Lathrop,
IIASA, Laxenburg, Austria

Stephen R. Watson,
Emmanuel College, Cambridge, U.K.

October 1981
CP-81-28

A version of this paper has been submitted for journal publication.

INTERNATIONAL INSTITUTE FOR APPLIED SYSTEMS ANALYSIS
2361 Laxenburg, Austria

ABSTRACT

The implementation of a nuclear waste management technology raises several issues concerning the regulation of social risk. The most basic of those issues are how to regulate a technology when the uncertainties in social consequences are important, and how to incorporate the relevant social values in the regulations. This paper presents a decision analytic approach to resolving these issues, based on the development of radiological risk evaluation indices. While it is essentially a case study, describing work carried out for the U.S. Nuclear Regulatory Commission, this case is used to discuss the more general issues involved.

We begin by discussing the need for risk evaluation to provide a clear and defensible basis for regulating technologies involving social risk. We then present an overview of the problem of evaluating the risk of nuclear waste management, and an overview of our approach. That is followed by a development of risk evaluation indices for the regulation of nuclear waste management. The indices developed are expected utilities, based on preferences elicited from groups of people. The use of the indices developed is illustrated in a hypothetical example. Finally, components of the analysis requiring further development are identified and discussed, and the usefulness of the methodology evaluated.

Running head: Evaluation of Nuclear Waste Risk.

DECISION ANALYSIS FOR THE EVALUATION OF RISK IN NUCLEAR WASTE MANAGEMENT

John W. Lathrop, IIASA, Laxenburg, Austria,
Stephen R. Watson, Emmanuel College, Cambridge, U.K.*

INTRODUCTION

Over the past two decades or so there has been increasing concern over the use of many modern technologies. While much of this concern results from the demonstrable and existing deleterious side effects of these technologies (for example, the effect on animal and plant life of pollution from many chemical plants, or the cumulative effects in the food chain of insecticides like DDT), there are some technologies for which the cause for concern is potential rather than actual (for example the depletion of the ozone layer by chlorofluorocarbons and by the exhaust gases of supersonic transports, or potential radiation releases from buried nuclear wastes). Governments are concerned to control both kinds of technological side-effects, but the appropriate regulatory action is much

* Authors are listed in alphabetical order.

more difficult to determine in the latter case, where there is considerable uncertainty about the possible nature of the consequences of the technology, and where they can occur in the distant future.

In this paper we discuss how the paradigm of decision analysis might be used to establish an index which can evaluate the risk associated with activities of the latter kind; in particular we shall exemplify our suggestions throughout by reference to the problem of regulating nuclear waste management. At the outset we ought to stress that we are not using 'risk' in the sense normally used in decision analysis. In bringing together decision analysis and risk evaluation, we must adopt terms that clearly distinguish concepts often confused between the two fields. In particular, 'risk' will be used to describe the potential for deleterious consequences associated with a technology, while 'uncertainty' will be used to describe the lack of information available concerning what the impacts of a technology might be.

The need for risks to be analyzed if they are to be adequately managed is beginning to be widely recognised. The important studies by Rowe[1] and Lowrance[2] on the nature and management of risk have been followed by a collection of papers outlining research needs and opportunities in this area, edited by Kates[3]. Our goal in this paper is to add to the emerging literature of this discipline by describing a particular approach to risk assessment, using decision analysis. This is a controversial area; for example, the criticisms by Lovins[4] and Cochran[5] of the decision analysis of alternatives for electricity generation carried out by Barrager et al[6], are paralleled by Hoos' attack[7] on the use of risk analysis for nuclear waste management and the critical comments of Tocher[8] concerning the use of decision analytic concepts in any social planning activity. These authors give many detailed arguments against particular quantitative methods; the reader is urged to consult them as an antidote to overconfidence in the methods of systems analysis, and decision analysis in particular. However, it is fair to say that most of their criticisms are destructive in nature and often might easily be paralleled by similar criticisms of alternative informal and traditional methods of decision making and regulation setting. It is our strong belief that if used properly, the quantitative methods of decision analysis, while subject to some valid criticism, improve decision making in that they provide a consistent base for analysis and improve communication. In many cases the informal methods share all the flaws of the formal approach; it is just that their very informality obscures the fact that these flaws exist.

In the next Section we give an overview of the problem and the approach we take. That is followed by the kernel of our paper, where we describe in detail a decision analytic approach to the evaluation of risk and how this method was used to construct risk evaluation indices for the regulation of nuclear waste management. The next section describes the results of our work in terms of a hypothetical example, which also serves to illustrate the role of the approach in the development of regulations. We then discuss some of the findings and the aspects of our approach which especially warrant further development. Finally, the last section briefly comments on the present and potential usefulness of the approach. This work was carried out for the U.S. Nuclear Regulatory Commission (NRC) under the direction of Lawrence Livermore Laboratory. Larger reports (Lathrop [9], Watson and Campbell[10]) contain more detailed descriptions of the project and the results of the study.

OVERVIEW: THE PROBLEM AND THE APPROACH

Before going on to discuss our work, we will briefly describe the nuclear waste management problem, characterize some special aspects involved in evaluating the risk of waste management strategies, and outline our approach.

Several activities involving nuclear material (electric power generation, weapons manufacture, medical treatment, etc.) produce radioactive waste. Some of the waste will remain highly radioactive for tens of thousands of years. A large stock of such wastes has already accumulated and more is being generated every day. In the U.S. the existing stock of highly radioactive waste is currently being held in temporary storage until a strategy for long-term disposal is established. The strategies currently being considered all involve the treatment and packaging of waste into forms designed to minimize leakage, then the burial of those forms in deep, geologically stable repositories. The transport, storage, and packaging of waste expose members of the current generation to radiation: low level routine exposures and a risk of accidental high-level exposures. In addition, even with the safest repository design there is a positive probability that radioactive isotopes will leak out and enter the water supplies and food chains of future generations, adding to their low-level radiation exposure. While there is a great deal of uncertainty concerning the effects of low-level radiation, such radiation could result in additional cancers, mutations, and cases of impaired fertility in future generations. In regulating nuclear waste management, then, the NRC controls both the overall level of risk and how that risk is distributed between current and future generations.

It can be seen from this description that there are four basic aspects that characterize the problem of evaluating waste management risks. First, nuclear waste management involves actions conducted now that effectively commit society to scenarios with probabilities of human health impacts over centuries. Thus many people may be affected that are not yet born. Second, the health impacts are best described on several different dimensions: there are several types of health effects, and they may be incurred by nuclear waste workers or people outside the industry in each of many generations. Third, a great deal of uncertainty is involved in assessing those impacts: in the ability of the waste management system to contain the wastes, in the effects of the extremely low-level radiation most apt to result from any release, and in the doses actually received (Who knows where people will live, in what densities, with what diets, in 10,000 years?). Fourth, while several agencies and companies are to be involved in nuclear waste management, risk management can be considered to be centered in one agency: the NRC.

The four aspects just listed characterize a situation where a quantitative approach to risk measurement would be valuable. In particular, the complexity of evaluating multidimensional and uncertain outcomes that affect people not yet born is a strong argument for the use of quantitative, as opposed to informal, risk measurement in support of risk management decisions. Quantitative approaches to risk measurement have quite a long history. (See Farmer[11], Starr[12] and the bibliography of Clark and Van Horn[13].) Broadly speaking, they have all recognized that a measure of risk should be an increasing function of the probabilities of deleterious consequences, and the severity of those consequences. A measure often chosen has been expected fatalities. Shortcomings of this criterion include its inability to cope with attitudes toward uncertainty (Who feels that a 50-50 chance of two deaths is just as bad as one certain fatality?), and the exclusion from consideration of consequences other than death (which in the case of radiation exposure includes effects as significant as genetic mutations). Papp et al[14] recognized this inadequacy and suggested that it should be rectified by using utility theory; the present paper describes how we have followed this suggestion for nuclear waste management.

While our decision analytic approach to risk measurement is described in more detail later, its essential concepts are presented in a brief overview now, both to orient the reader for the rest of the discussion and to provide a framework for defending this approach against the main points of its critics. Very briefly put, the approach consists of developing a multidimensional utility function over the health-effect consequences of a nuclear waste management system, separately assessing a probability distribution over those consequences for each alternative system, and then calculating the risk index for each system as its expected utility. Because our utility function increases with the severity of negative consequences (contrary to convention), the risk index is an increasing function of probability and severity of consequences, as desired. As explained by Howard,[15] the expected utility represents preferences for uncertain outcomes on a cardinal scale, ranking complex alternatives in a manner consistent with preferences revealed in comparisons of simple alternatives.

We adopted the particular form of multidimensional utility function developed by Howard.[15] That function consists of a multidimensional value function to represent value trade-offs between the various dimensions under certainty, then a von Neumann-Morgenstern utility function on that value function to represent attitude toward uncertainty. In classic decision analysis applications, the function is fitted to the preferences of the decision maker by means of an elicitation protocol. The protocol consists of a series of simple paired comparisons (Which do you prefer, A or B?) designed to reveal the decision maker's value trade-offs and attitude toward uncertainty. The essential idea of the approach is to fit the utility function to simple comparisons, e.g., where the alternatives differ on only two dimensions, where the decision maker's responses are apt to reflect his actual preferences. The utility function can then be used to make very difficult comparisons among alternatives with complex and uncertain outcomes in a manner consistent with the responses to the simple comparisons. In this application, however, there was no single person whose values were to be represented. Since the NRC was in some sense the decision maker, we decided to develop a methodology to represent the values of four groups of people the NRC could choose to be responsive to. Those groups are described later. Rather than attempt to aggregate values over the four groups, values were aggregated into a single utility function within each group, so that four separate risk evaluation indices were developed.

Now that the broad outlines of our approach have been presented, it is useful to examine them in light of the chief criticisms often raised against utility approaches. A useful compendium of criticisms is found in a paper by Tocher[8], where he raises three key points often raised by others. His first criticism is that utility methods force N-dimensional alternatives onto a single-dimension scale via value trade-offs elicited by posing hypothetical choices to individuals. Our answer to this point is simply that any method to evaluate the risk of waste management systems must have the same shortcomings, though such problems are apt to be concealed in the unexplicated parts of less formal methods. Any comparison of the risks of alternative systems must effectively project those risks onto some single, common scale if anything is to be said along the lines of: System A presents more risk than System B. Granted, there are problems with any use of index numbers to evaluate multi-dimensional phenomena, whether those indices are IQ, GNP, or risk. Information must be lost on the way from N dimensions to one dimension. Yet the fact remains that the NRC is charged with managing something called risk. It is our conviction that effective risk management requires some single-dimension definition of risk readily understandable to the NRC, the people it is to regulate, and the people it is to protect. The value trade-offs within that definition must be hypothetical, because decisions concerning projected waste management systems involve choices among sets of effects that have never occurred before.

Tocher's second criticism is that the axioms of utility theory serve to represent individuals as mechanistic beings without free will. This point is only peripherally related to the waste management case, since utility theory is used in this case to represent preference, not predict behavior. Clearly, some determinism of preference must be assumed if individuals' values are to be incorporated into choices among current actions with long-term impacts. Tocher's third point concerns the difficulty of aggregating values. This difficulty was mitigated in the methodology reported here by limiting aggregation to within groups where some homogeneity of values was expected. In delivering four different risk evaluation indices to the NRC, based on the values of four disparate groups, the critical value aggregation problem was left to the political processes of the NRC.

DEVELOPMENT OF A RISK EVALUATION INDEX

While the basic ideas of the decision analytic approach adopted here are relatively straightforward, there are several steps in its implementation. We shall now describe those steps as we encountered them in our development of a risk index, some of the problems each step entailed, and how we dealt with them.

Consequence Scope

In constructing a utility function describing consequences, we first had to specify the important possible consequences of a nuclear waste repository. The most significant possible effects are, clearly, health effects on humans, but there are many others, such as effects on animal and plant life, restrictions of civil liberties and restrictions on land use. We made a modeling decision at this stage to limit attention to health effects on humans. This limits, of course, the use to which our risk index can be put, since in comparing repositories any of the other possible consequences may be of significance. However, the purpose of this study was not to analyze specific decisions on the location and design of repositories, let alone whether or not radioactive waste should be generated; it had the lesser, but still important goal of providing a measuring device for only one aspect of the waste management problem, namely the health risk to humans. Here is an example of the methods of decision analysis contributing to public policy formation by analyzing part of the problem. In our experience, such partial analyses are much more likely to be used and useful than any attempt to bring the whole decision making process for an issue of public policy under the hammer of hard analytic methods.

Value Source

The idea of using decision analysis on public policy issues is, of course, not new. Both Howard[15] and Edwards[16] have proposed schemes for social decision analysis. One of the problems which such schemes raise, however, is the determination of whose values should be represented in the utility function. In an evaluation of the risk in waste management it would seem desirable somehow to reflect public values; but these may differ importantly from person to person--how may such differences be combined into one index? It should be helpful here to find some method of aggregating individual utility functions to create a group utility function. While there are difficulties in specifying such a procedure (as Seaver[17], p 14, observes, "no entirely satisfactory method for devising group utilities exists"), the work of Keeney and Kirkwood[18] cites theoretical support for the idea of using a weighted sum of utility functions as a group utility function, the weights to be determined by a "benevolent dictator" or an "honest broker". At the same time, the weights given to people with markedly different stands on a policy issue represent trade-offs perhaps best made in the course of the political process, not in the risk analysis.

The relevant political process in this case centers on the NRC, which is effectively charged with balancing the political interests of any conflicting parties. There is always reason to suspect that any regulatory agency is subject to disproportionate pressure from the industry it is supposed to monitor, so that such an agency may not be an ideal umpire of conflicting interests (see, e.g., Hoos[7]). However, the goal of the approach described here is not to reform the political process, but simply to develop a risk evaluation methodology that incorporates social values into the existing political process. Consequently, we drew the bounds between the evaluation and the process it is to serve, and decided to represent separately the values for four groups, described below. The groups chosen were ones to which the NRC is responsive, which span the political spectrum of interests faced by the NRC, and within which values could be expected to be relatively homogenous. While this strategy avoids aggregating values across conflicting groups of people, values must still be aggregated across individuals within each homogenous group to arrive at a group utility function. That aggregation is explained later.

In the study, 58 people were interviewed and utility functions elicited from each of them. They were divided into groups as follows:

a. National Advisors (13 respondents)

   This group consisted of persons who had the ear of the Federal policy-makers in that they either served on nuclear waste advisory bodies, or their views were published or otherwise consulted by Government policy-makers.

b. Concerned Citizens (33 respondents)

The intent in this case was to select citizens who were at least somewhat abreast of social issues and who were concerned that Government actions should reflect the general public interest.

c. Nuclear Power Opponents (7 respondents)

These were persons who were known publicly to oppose further development of nuclear power, at least until safety problems have been resolved.

d. Nuclear Power Advocates (5 respondents)

These were persons who had been identified as advocates of the further development of nuclear power. Some of the respondents in this group maintained that they did not consider themselves as advocates of nuclear power.

Group utility functions for each of these four groups were calculated as weighted sums of the individual utility functions. Then these were combined with example probability distributions describing different repositories, producing four distinct risk evaluation indices for each repository. We return to the results of these calculations later. Because of our decision to elicit utility functions from so many people, it was necessary to construct a simple standard form for the utility functions, which called for several approximations in the following steps.

Consequences

One of the basic problems of approximation concerned how to describe health effects. The first widely used and definitive discussion on the health effects of radiation is the 1972 report of the U.S. National Academy of Sciences Advisory Committee on the Biological Effects of Ionizing Radiation[19] (the BEIR report). It is apparent from this and other documents (see ICRP #26[20]) that the health consequences of radiation exposure are many and varied. Despite this, there is a fairly obvious categorization, as portrayed in Figure 1. The first distinction is between 'stochastic' and 'non-stochastic' effects. To quote ICRP #26 [20], p.2:

> "Stochastic effects are those for which the probability of an effect occurring, rather than its severity, is regarded as a function of dose, without threshold. Non-stochastic effects are those for which the severity of the effect varies with the dose, and for which a threshold may therefore occur."

Stochastic effects can be further categorized as either somatic, if they become manifest in the exposed individual himself, or genetic, if they affect his descendants. It is clear that within each category there is a very large number of possible effects. The modeling decision that now faced us was how many of these to include in a list of attributes for a utility function. Here a balance had to be drawn, as always in applied decision analysis, between analytic simplicity and completeness. We needed to find a set of attributes which included everything important in the evaluation of risk, and yet was small enough for a utility function over the attributes to be elicited reasonably easily from a large number of people in a

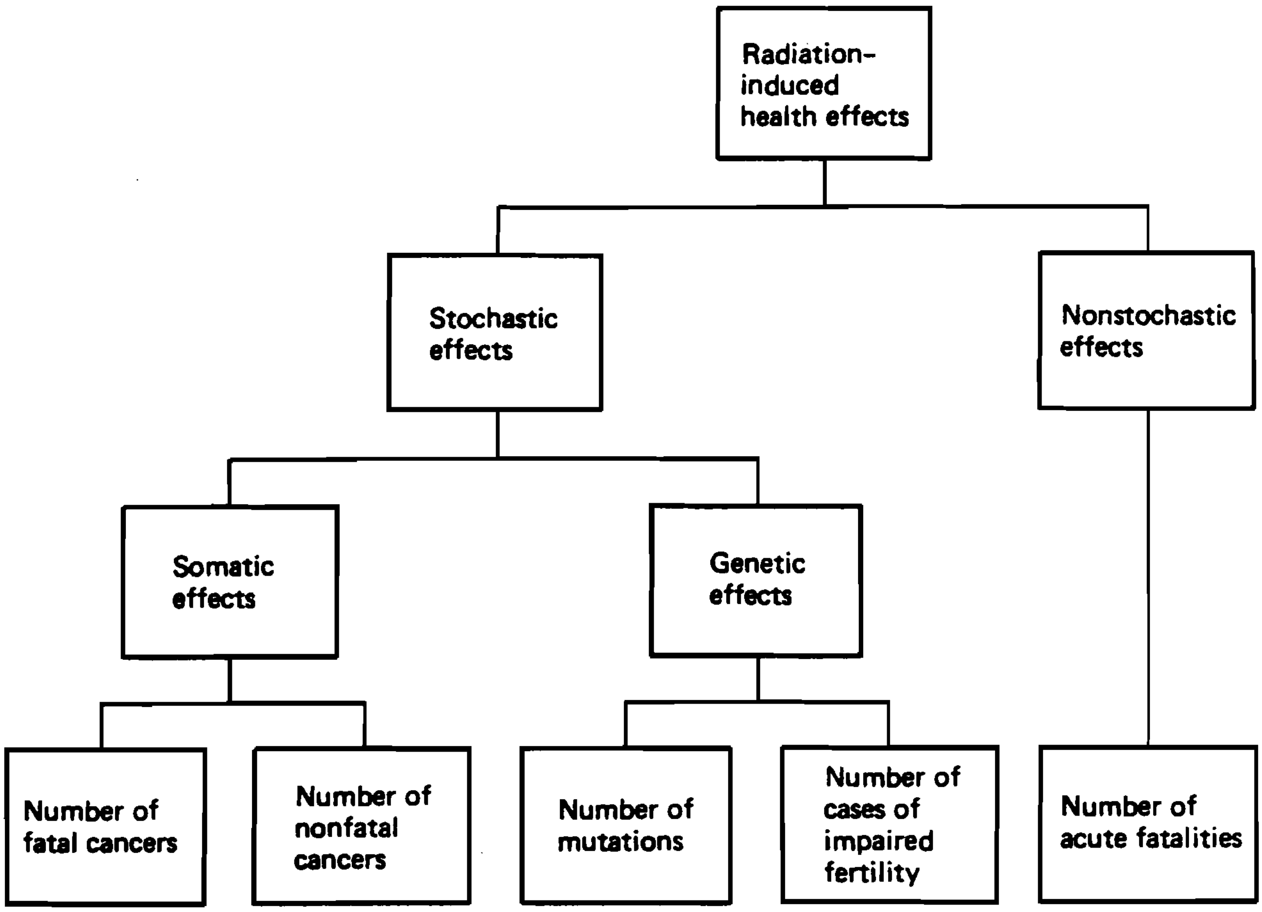

Figure 1. Categorization of health effects.

short time. We selected the following five variables, as indicated in Figure 1.

x_1 - Number of fatal cancers

x_2 - Number of nonfatal cancers

x_3 - Number of mutations

x_4 - Number of acute fatalities

x_5 - Number of cases of impaired fertility

Acute fatalities here refers to those that occur shortly after exposure, and so correspond to incidents of very high dose associated with pre-seal accidents.

Several approximations and assumptions about the respondents' value structure are implied by the above list. First, the use of the total number of cases of a particular kind implies that any such case is as bad as any other such case. Second, there are other known radiation effects which do not appear on the list, such as cataracts or retarded development in children. Third, some of these categories cover a very wide class and in order to enable trade-offs between such effects to be accessible psychologically to our respondents, it was necessary to particularize them considerably; thus the class of mutations was represented by a mentally sub-normal person, who needs some extra care throughout his life. This last approximation will only be good if this case is an approximate certainty equivalent of the class of all possible mutations; this in turn depends upon the nature of the probability distributions with which the utility function will be used, and we had to complete this study without a good idea of those probabilities.

We are not satisfied that this list of attributes is the best set to describe health effects, and we feel that more work should be done on determining a good set. However, this seems to us a reasonable first attempt, and does satisfy some of the criteria for choosing attributes: sufficient complexity to cover the important structure of the problem, yet sufficient simplicity for elicitations of the trade-offs to be believable and available in a reasonably short period of time.

The next point concerns the circumstances in which the radiation doses arise. As Fischhoff et al[21] discovered, the circumstances of risks do affect their importance as judged by members of the public. Of the many possible circumstances, the two we judged to be significant enough to include explicitly in the utility function were whether or not the risk was undertaken as part of a person's occupation, and the time at which the radiation dose might arise. Figure 2 illustrates this categorization. Notice that occupational risk can only be suffered by workers in the current generation, since after sealing no further operation of the repository will be necessary.

We define an index i to indicate the circumstances of the dose giving rise to a particular health effect according to the following scheme:

$i = 1$: Effects due to occupational exposure before sealing.

$i = 2$: Effects due to nonoccupational exposure, to the current generation.

$i = 3$: Effects caused by a dose in the two thousand years following repository sealing.

$i = 4$: Effects caused by a dose more than two thousand years following repository sealing.

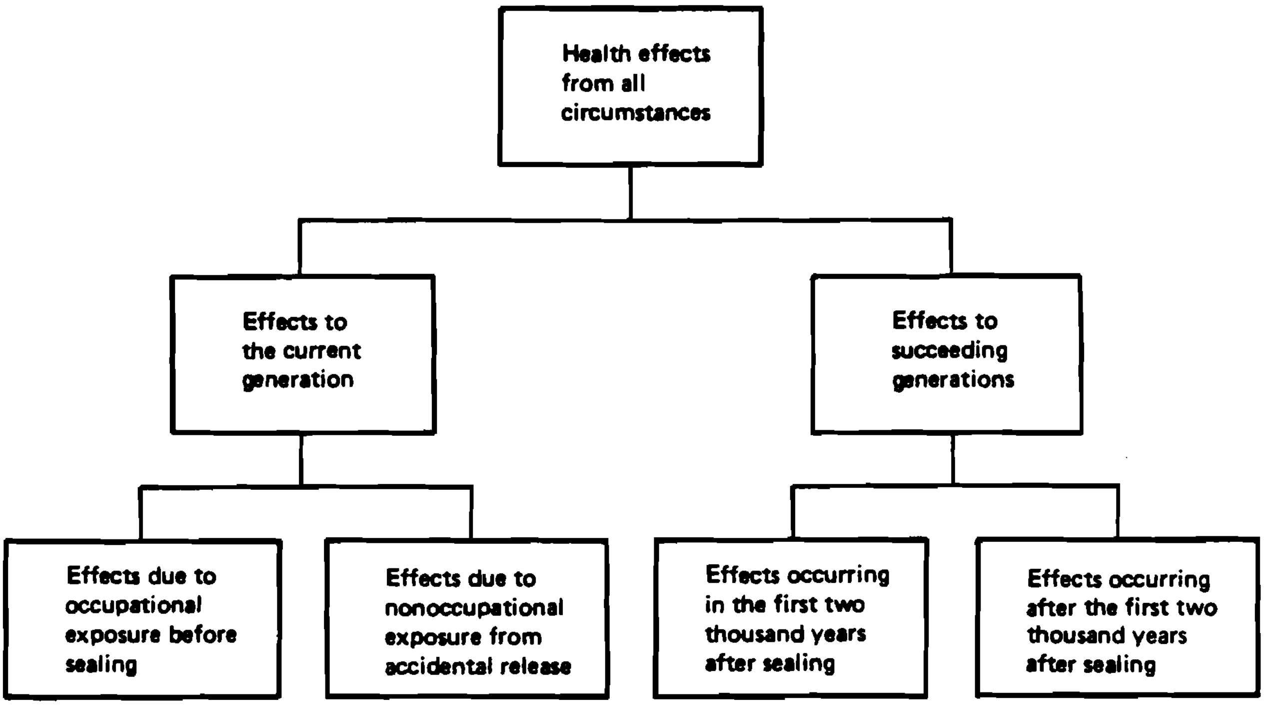

Figure 2. Categorization of the circumstances of radiation exposure.

Each of the five health effects can occur as a result of circumstances in each of the four categories above. There are, therefore, twenty variables x_{ij} which describe the possible effects caused by a nuclear waste repository, where $i = 1,\ldots,4$ indicates the cause of the dose and $j = 1,\ldots,5$ indexes the type of health effects (e.g., x_{32} is the total number of nonfatal cancers arising from doses received from the nuclear waste in the two thousand years after the repository has been sealed). Note that the time division here is somewhat coarse. Once again, a balance had to be struck between capturing the essential elements of the respondents' value structures and keeping the model reasonably simple.

Individual Utility Functions

The next problem was to establish a parametric structure for the utility function of each individual. We chose to construct first a value function $v(\mathbf{x})$ (where $\mathbf{x}$ is the vector whose twenty components are x_{ij}), which would have the property that any possible set of health effects $\mathbf{x}$ judged to be equivalent to or worse than any other $\mathbf{y}$ would have $v(\mathbf{x}) \geq v(\mathbf{y})$. The form we adopted was

$$v(\mathbf{x}) = \sum_{i=1}^{4} \sum_{j=1}^{5} \alpha_i \beta_j x_{ij}$$

This form makes some rather sweeping assumptions about preference structures. First, the linearity implies that, for example, if one extra acute fatality is as bad as two extra mutations at one level of all the variables, the same is true at all other levels of the variables. Second, the fact that the constant coefficient of each variable x_{ij} is a product $\alpha_i \beta_j$ implies that trade-offs among health effects do not depend on the circumstances of the dose, and trade-offs between the same effects in different circumstances do not depend on which health effect is considered (i.e., if three acute fatalities in a thousand years time is as bad as two acute fatalities now, then three fatal cancers in a thousand years time is as bad as two fatal cancers now). Although these assumptions are fairly strong, our respondents' preference structures seemed to be consistent with them to an adequate approximation.

The final stage in the construction of a von Neumann-Morgenstern utility function for each of our respondents was to reflect uncertainty preference (our synonym for the more usual 'risk preference,' which we shall avoid to prevent confusion with the more general use of the word 'risk' in this paper). We constructed a single-argument utility function reflecting uncertainty preference, whose argument was $v(\mathbf{x})$. In our detailed reports[9,10] we used an exponential family of utility functions, which has the advantage of being a single- parameter family. This proved to be an inadequately rich family to describe the uncertainty preferences of many of our respondents, at least as they were assessed in this study. In a re-analysis of our work, a two-parameter family of utility functions would be desirable.

Note that this utility function incorporates a very straightforward solution to a central problem with risk evaluation of nuclear waste management: how much relative weight to give health effects to future generations as opposed to health effects to the current generation. While this characterization of the problem puts it in the obscure language of evaluation, it can be given much more meaning in terms of a decision dilemma: for the money remaining in your limited budget, you can save (expectationally) either 10 lives in the current generation, or 11 lives 1,000 years from now. Which do you choose? There are a set of arguments for discounting future health effects, reviewed in Goodin[22], which would lead one to prefer saving 10 current lives over 11 future lives. However, as Goodin points out, the economic opportunity-cost arguments for discounting do not readily apply to the long time periods involved in nuclear waste management, since an appropriate rate of return on resources saved is extremely difficult to estimate, even in the form of a subjective probability distribution. Goodin goes on to describe and attack various other rationales for discounting. The strength of our approach is that it does not presume any particular prescriptive rationale. It simply asks respondents, in the trade-off format explained above, for their relative weights between current and future health effects. A respondent is free to take a stand that future health effects should be weighted equally to current ones, or to ascribe to a discounting rationale and set his or her own rate of discount (positive or negative). If there were a clear consensus from a broad community of experts as to some correct, non-intuitive approach to relative evaluations of health effects over time, our value elicitation could be attacked. But clearly, as Goodin's article attests, there is no such consensus, so it could be argued that our respondent's answers to this particular trade-off question represent as good a basis for setting policy as conflicting expert opinions.

Value Elicitation

Having developed the structure of the utility function, the values of the parameters had to be elicited from each individual respondent. One of our team sought out and interviewed 58 respondents (as described above) and made estimates of the parameters $\{\alpha\}$ and $\{\beta\}$ for each person interviewed. He did this by first asking the respondent to order the five health effects as to seriousness, and then making pairwise comparisons. For example, if the respondent stated that a fatal cancer was worse than an acute fatality, he was asked how many acute fatalities was as bad as one fatal cancer. By such comparisons and associated cross-checks the parameters $\{\beta\}$ were estimated. Similar comparisons, only this time for a particular health effect, but with different underlying causes, gave rise to estimates for the parameters $\{\alpha\}$. Finally, lottery type questions were asked of the respondents, from which their uncertainty preference was assessed.

Because of time limitations it was only possible to spend an hour or so with each respondent. Given the known effects of expressed value changing over time and depending on the type of questions asked (the problem of labile values, Fischhoff et al[23]), we have considerable doubts as to whether our results would have been the same if a longer time had been taken in eliciting values.

Value Aggregation Within Groups

The next step in the calculation of risk indices for each of the four groups mentioned above was the determination of the weights to apply to the utility function of each individual within each group to obtain the group utility function as a weighted sum. As mentioned before, according to Keeney and Kirkwood[18], and Keeney and Raiffa[24], p.539, one reasonable approach to determine the weights in an additive group utility function is to use the judgments of a "benevolent dictator" or "supra decision-maker," who makes a fair balance between the intensities of preference of the group members. This would seem a particularly appropriate approach in this case, since the U.S. Nuclear Regulatory Commission has the responsibility of regulating the nuclear industry in the U.S. in the balance of the public interest. One could ask the Commissioners themselves to determine relative weights between individuals within each group. The nature of our study, however, precluded us approaching the Commissioners. We had to resort, therefore, to a procedure for determining weights based on the concept of equal weighting of calibrated utility functions.

The process of calibrating individual utility functions consists of weighting each function so that a utile on one is in some sense comparable to or commensurate with a utile on another. Even this limited operation involves interpersonal comparisons of utility, and so has no fully satisfactory methodology. Our approach was to standardize the utility functions on a unit interval: one fatal cancer caused by a pre-seal nonoccupational low dose. This standard was chosen because it was the most clearly understood health effect, so that the effect of differing individual interpretations of a health effect on calibration error was minimized. In addition, this standard involved the dose circumstance generally used as a basis by the respondents against which other dose circumstances were discounted, so that the effect of differing discounts on calibration error was minimized.

Probability Distributions

Now that we had constructed four utility functions over the twenty-component health effects vector $\mathbf{x}$, each derived from different groups of people, we could combine these utility functions with probability distributions over $\mathbf{x}$ describing the uncertainty in the consequences from any particular repository. Here we faced a further problem. Although the theory of subjective probability implies that probability distributions can be constructed describing any degree of uncertainty, the implication of the methods of social decision analysis suggested by Howard[15] and Edwards[16] is that the best available expertise should be harnessed to construct probability distributions on the outcomes of alternative repositories. Lawrence Livermore Laboratory was at the time of this study engaged in an extensive project to produce the necessary expert knowledge of the likely consequences of any particular repository site and design, but the appropriate results were not available at the time our reports[9,10] were completed. Instead, in order to test the implications of our utility functions, we elicited example subjective probability distributions from two staff members at Lawrence Livermore Laboratory who had considerable personal experience in the assessment of radiation hazards, and in the possible consequences of geological disposal of radioactive waste. They produced probability distributions for two hypothetical repositories, differing from each other in a way that was meant to represent the effects of a hypothetical regulation requiring an increased amount of waste packaging. It was hypothesized that the regulation would increase the probability of pre-seal health effects slightly (due to increased handling), and decrease the probability of post-seal health effects (due to increased isolation). It should be emphasized that these probability distributions are for example purposes only, and are in no way based on the results of the physical repository modeling effort conducted at Lawrence Livermore Laboratory.

RESULTS

The results of our work fall into two categories: features of the elicited values and characteristics of the risk evaluation index itself, as generated from the assessed preferences of the four respondent groups and applied to the example probability distributions mentioned above.

Elicited Values

The two most interesting features of the elicited values happen to coincide with the two main improvements of expected utility over expected fatalities as a risk measure: the representation of attitudes toward uncertainty and the evaluation of more dimensions of health impacts than fatalities alone. Concerning attitudes toward uncertainty, half of the respondents were uncertainty preferring (i.e., risk seeking) in that they preferred a one percent chance of 100 fatal cancers over one fatal cancer for certain. About one quarter of the respondents were uncertainty averse, one quarter uncertainty neutral. In every one of the panels, taken separately, less than half of the panel members were uncertainty averse. These preferences indicate that a group utility function that assumes individual uncertainty attitudes to be neutral would match or overstate the aversion to uncertainty of about three-quarters of the respondents. Because the risk evaluation index is intended to measure the risk of repositories relative to the risk of standards involving less uncertainty, and because any errors in the estimates of repository impacts are apt to understate their uncertainty (see Goodin[22]), an index based on a group utility function that overstates aversion to uncertainty could be considered desirably conservative. For these reasons, and because of several advantages for implementation of an uncertainty neutral group utility function, that form of function was adopted for the risk index calculations performed in this study.

Concerning evaluation of health impacts other than fatalities, value trade-offs between fatal and nonfatal health effects were such that nonfatal health effects contributed significantly to the risk evaluation index. As one example, more than half of the respondents considered a mutation as worse than or equivalent to a fatal cancer.

Characteristics of the Index

The clearest way to present the risk index, and the implications for that index of the elicited values, is to demonstrate its use. The following paragraphs will step through a few simple example calculations to that end. The calculations are oriented toward answering three basic questions: does the proposed regulation decrease overall risk, by how much is the overall risk decreased, and does the regulation decrease the risk to an acceptable level.

The most elemental question these risk evaluation indices can help to answer is: Does the proposed regulation decrease overall risk? For a regulation requiring increased packaging of wastes, for example, the answer is not immediately clear, since it would decrease post-seal risk at some expense in increased pre-seal risk. The risk indices (expected utilities) calculated according to the scheme outlined in the previous section are given in Table I, for repositories with and without the regulation, for each of the four respondent groups. Since we adopted an uncertainty neutral (linear) form for the group utility functions, and because of the particular normalizations we used, the numbers in Table I are not only expected utilities; they are also equivalent pre-seal nonoccupational fatal cancers. As the column differences in Table I make clear, the example regulation does in fact increase overall safety for each respondent group.

TABLE I

Expected Utilities of Four Respondent Groups for two Repositories

Repository	National Advisors	Concerned Citizens	Nuclear Opponents	Nuclear Advocates
repository A, without new regulation	2.4	9.1	5.2	.46
repository B, with new regulation	2.0	7.3	4.2	.41

Differences in index value between groups can be explained in terms of preference differences in discounting of future health effects and in the relative importance of different health effects. However, there is no direct operational significance of the differences within rows in Table I. All that matters as far as regulation selection is concerned are the differences within columns. This is because utility measures only have significance for actual choices faced by some decision making process, based on a single assessed set of values. The actual choice faced here is between repositories A and B, using the values of each of the four groups, respectively. A comparison of, say, the National Advisors' utility for A and the Concerned Citizens' utility for A is meaningless: there is no actual choice between A and A, and there is no single set of values involved. A discussion below compares the utilities of the different groups in a more meaningful way by presenting implications of differences in utilities for determinations of acceptable risk.

There is one basic concept perhaps not stressed enough in this report that Table I helps make clear. The differences in numbers within either row of the Table should not be interpreted as measurement error. We have not constructed a single risk evaluation index, but a set of four such indices, each capturing, in some way, the attitudes of one of the four panels. The risk evaluation index of a repository is not some physical attribute of that repository; rather, it is an evaluation of the impacts of that repository, as judged from a particular set of personal values.

The second basic question the risk evaluation index can help to answer is: How effective is the proposed regulation? Table II presents three different ways to scale the difference in risk evaluated between repository A (without the regulation) and repository B (with the regulation). These differences are measures of the effectiveness of the proposed regulation. The first row of Table II is simply the set of differences between rows from Table I: the reduction of risk in utiles caused by the regulation. Because of the scaling used, those utiles are equivalent pre-seal nonoccupational fatal cancers prevented by the regulation, and so are measures of the equivalent lives saved by the regulation, for a very particular type of life saving. The utility functions used here are linear, so each function is basically a set of coefficients between arguments. It is an easy matter, then, to convert any difference in utiles (as in Row 1) into any set of utility function arguments (health effects), and then convert to the dose that would produce that set of health effects, using the dose effectiveness coefficients of the BEIR Report[19] used throughout this study. Two sets of results of such conversions are displayed in Rows 2 and 3 of Table II, equivalent reductions in occupational dose and pre-seal nonoccupational dose, respectively. The logic of the utility function is such that for any column in Table II the dose reduction in Row 2 would cause the risk evaluation index reduction in Row 1, as would the dose reduction in Row 3. Another way to put it is for any group (column in Table II), the dose reduction in Row 2 would be worth just as much as the dose reduction in Row 3 or the reduction in the number of fatal cancers in Row 1. Of course, the utility functions differ between groups, so the ratios between entries in Table II vary from column to column.

TABLE II

Measures of Risk Reduction due to the Proposed Regulation

Measure of Risk Reduction, Difference in Equivalent:	National Advisors	Concerned Citizens	Nuclear Opponents	Nuclear Advocates
...pre-seal non-occupational fatal cancers	.44	1.8	.96	.05
...occupational dose (man-rem)	12,000	34,000	14,000	3,000
...pre-seal nonoccupational dose (man-rem)	5,300	19,000	8,800	800

The second and third rows of Table II can be useful for comparing the effectiveness of a regulation with other regulations or technical alternatives on convenient dimensions. For example, if an alternative regulation to the one used in the example would have the sole effect of reducing occupational dose, Row 2 would offer a very direct comparative measure.

The effectiveness of the regulation could be measured in terms of reduction in expected fatalities, a very different measure from a reduction in equivalent fatal cancers. For the hypothetical probability distributions used in this example, the reduction in expected fatalities due to the regulation comes to .23 expected lives saved. It is interesting to compare this figure with the first row of Table II. For every group except Nuclear Advocates, the fact that the reduction in equivalent fatal cancers incorporates nonfatal health effects leads to measures of regulation effectiveness more than twice as large as the less comprehensive "expected lives saved" measure. On the other hand, the fact that the Table II - Row 1 measure incorporates the Nuclear Advocates' discount factors for future and occupational fatalities leads to a measure of regulation effectiveness much smaller than the undiscounted "expected lives saved" measure. These examples should make clear that the "expected lives saved" measure is not at all "value-free." It in fact makes specific assumptions concerning value trade-offs that seem to be importantly at variance with the value trade-offs assessed from our respondent groups.

The third basic question the risk evaluation index can help to answer is: Does the regulation reduce the risk of the repository to an acceptable level? The definition of acceptable risk is quite involved and will not be addressed here (see Lathrop[9] and Watson and Campbell[10]). But whatever the definition of acceptable risk, the risk evaluation index can play a key role in its determination by providing a common scale on which to compare risks, and on which to set an acceptable risk limit. For example, suppose some analysis finds that a repository is acceptable if its risk is less than the risk due to an occupational population dose of 65,000 man-rem (maximum individual dose less than 5 rem). For any given risk evaluation index, this dose can be expressed in utiles as a limit on that index scale. Any repository can be evaluated using the risk evaluation index on the same utile scale, and so directly compared with the acceptable risk limit. Of course the risk evaluation index depends on the values represented, so such an acceptable risk limit provides yet another way to explore the implications of differences in values between our four groups.

In our case the risk evaluation index is simply a weighted sum of four expected doses (occupational, nonoccupational in each of three periods), with different weights for different groups. Those differences in weights mean that any given acceptable risk limit may lead to different relative positions of the limit and the evaluated repositories for the different groups. For example the 65,000 man-rem limit mentioned above was selected solely because it leads to corresponding index values for the National Advisors and Nuclear Opponents groups such that repository A is unacceptable and repository B is acceptable. In other words, based on the 65,000 man-rem limit and the values of those two groups, the proposed regulation reduces the risk from an unacceptable to an acceptable level. However, the Nuclear Advocates discount the future health effects of the repositories so much that the risk evaluation indices of repositories A and B both come out below the index value of the same 65,000 man-rem current dose limit. At the same time the Concerned Citizens weight future health effects so highly that their risk evaluation indicies for both repositories come out above the index value of the 65,000 man-rem current dose limit. That is, using the same basis for an acceptable risk limit, the Nuclear Advocates' values are such that both repositories are acceptable, while the Concerned Citizens' values are such that both repositories are unacceptable. Lathrop[9] describes these calculations in more detail, and presents other example acceptable risk limits, including one based on comparison of the repository risk with risk allowed by an existing regulation.

DISCUSSION

The purpose of this paper has been to describe an attempt to use the techniques of decision analysis to establish risk indices for nuclear waste repositories which reflect public values. The novel aspect of our approach, which we have not noticed reported elsewhere, has been the creation of utility functions representing segments of public opinion by first constructing utility functions for a sizeable number of individual respondents and then combining them to form group utility functions. One problem of this approach was that difficulties in eliciting values from so many respondents required us to make rather more approximations in the methodology than would normally be necessary in a study of this kind.

The two primary elicitation results call for some discussion. First, the small fraction of uncertainty averse respondents suggests that if a risk evaluation index is to reflect the popularly observed aversion to catastrophe, it must represent that feature in some other way than the uncertainty aversion of utility theory. Second, the large weight given to nonfatal health effects in our study demonstrates the importance of a risk evaluation index comprehensive enough to include them.

Several interesting results came out of this research and are presented above. However, the most valuable results are the three very general aspects of risk evaluation that our approach has made clear by example. First, risk is not some physical quality of a physical system; it is a function of both the physical system and the group of people evaluating the risk. Second, there is no value-free measure of risk; even the commonly used expected fatalities measure assumes particular values.

Third, while our methodology defines a scale on which to set an acceptable risk limit, that scale is so value-dependent that the relative positions of repositories and the acceptable risk limit can change as a function of the group whose values are represented.

To lay the foundation for our conclusions, it would be helpful now to review and summarize the main aspects of our approach that require further development.

a. Group utility functions

We have already mentioned that there is no firmly established methodology for deriving a group utility function by some combination of individual utility functions. Even if we allow the notion of a group utility function because of its intuitive reasonableness, a combination of individual utility functions is not the only way to achieve it; we could, for example, determine the utility function by group discussion and argument to produce unanimity about its structure. Despite this, the theoretical background discussed by Keeney and Kirkwood[18] supports our approach as a reasonable one. The actual determination of the weights that we have used here to calculate the group utility functions leaves a lot to be desired, however. Were the NRC to use this approach, more research would be needed to determine more defensible weights to apply to the different individuals in each group.

b. Structure of utility functions

There were several assumptions made in the construction of the individual utility functions that should be more thoroughly tested and, if necessary, relaxed, in a more extensive revision of this study. First we assumed that our restricted list of twenty health effects would be a rich enough basis for the utility function adequately to capture the relevant public values. Second, we assumed that the trade-offs between effects were constant, and, moreover, that trade-offs among consequences of a dose did not depend upon the circumstances in which the dose arose, and vice versa.

c. *Elicitation methodology*

Because of the large number of respondents from whom we elicited utility functions, it was necessary to spend only an hour or so with each person. It would have been desirable to spend much longer (say two days) with each of our subjects, but this was impossible not only because of the limited time available to us, but also because our respondents were busy people themselves. Given the difficulty of the issues we were asking people to address (who has thought about how much worse it is for a cancer to occur in the next 30 years than 10,000 years from now?), and the recognized difficulty in 'Knowing What You Want'[23], the utility functions we assessed should be regarded as first approximations only.

d. *Probability modeling*

It is vital for the application of this approach to risk evaluation that the probabilities should be sound in some intuitive sense. The probabilities used in the example presented above were for example purposes only, and require much improvement. The construction of probabilities in this area is exceedingly difficult and contentious. Even in the much 'harder' area of reactor risk analysis, where the final probabilities of failure depend to a large extent on known system design and failure rates of individual system components, for which extensive failure data is known, an attempt to determine the probabilities of failure met with contention (see the Rasmussen report[25], and its various critiques[26,27]). However, much of the criticism of that attempt was that it appeared to give an objective answer to the question of how likely a reactor accident would be. It is our view that probability is only a measure of degree of belief, and it is clear that in determining probabilities of radioactive release in tens of thousands of years time there must be a large measure of subjectivity, albeit based on available expert knowledge. We believe that it is possible to obtain adequate probabilistic descriptions of the consequences of a radioactive waste repository, but the necessary combination of the judgments of different experts (in material science, geology, demography and biology) has yet to be made.

CONCLUSION

In light of the further development required in the areas of our methodology outlined above, we must pose the question: Can we recommend that the NRC use the risk evaluation methodology described here to form a basis for the regulation of nuclear waste management, or are the problems so great that this whole approach to risk evaluation should be abandoned? Our answer is a positive recommendation to apply our methodology, but only after further development. It is clear, because of the difficulty of the problems outlined above, that considerable further work needs to be done before we can be confident that the risk evaluation indices created adequately represent the relevant social values, and provide a satisfactory means for incorporating those values into the risk management process. As challenging as those problems are, the fact remains that risk must be evaluated before it can be managed, and we feel that the risk evaluation methodology described here is an improvement over the less explicit process of risk evaluation that would probably be used in its stead. While the value elicitation questions called for in our approach were very difficult to answer, the fact remains that those questions must be answered in the course of nuclear waste risk management, either explicitly with a methodology such as the one described here, or implicitly without any formal analysis. All our methodology does is force people to confront these difficult trade-offs directly, rather than leave them to be determined implicitly by a process that manages risk without defining it.

Note that the methodology presented here is not intended to depoliticize what is clearly a political process. The determination of an acceptable level of risk is left entirely to the political domain. All that is suggested here is a measuring rod, so that the political debate surrounding the regulation of social risk is at least clarified to the extent that there is a clearer idea of what risk is. Even more than that, the concepts of risk proposed here are not normative in nature, but are based on social values elicited from groups of people to which the regulatory agency is normally responsive. As several people have pointed out (see for example Otway et al[28]) other issues, such as centralization of power, may be at stake in political debates ostensibly concerned with technological risk. But if risk is more clearly defined, then those other issues will be brought more clearly into focus, instead of being obscured in ill-defined notions of social risk.

On the basic argument that an improvement in the present level of information about public values concerning risk is necessary for the proper regulation of nuclear waste management, we maintain that our approach is worth pursuing. At the very least, it will provide a rational basis for proposed regulations put forward by the staff of the NRC at the start of the long chain of review. But beyond that, we believe that as society presses more toward democratic involvement in the regulatory process, the methodology presented here will become more and more attractive as a fair and just means to reflect social values in regulatory decisionmaking.

ACKNOWLEDGEMENTS

The work on which this paper is based was undertaken while Dr. Watson was on leave of absence with Decisions and Designs, Inc., McLean, Virginia, and Dr. Lathrop was with the Lawrence Livermore Laboratory, Ca., and was supported by subcontract 9693603 from Lawrence Livermore Laboratory. The authors wish to acknowledge the valuable assistance of Vince Campell, of Decisions and Designs, Inc. Dr. Campbell performed the value elicitations described in this paper, and made several helpful comments concerning earlier drafts of the paper. The authors also wish to acknowledge several helpful criticisms by referees on an earlier draft.

REFERENCES

[1]W.D. Rowe, *An Anatomy of Risk*, New York, N.Y., Wiley, 1977.

[2]W.W. Lowrance, *Of Acceptable Risk: Science and the Determination of Safety*, Los Altos, Ca., W. Kaufmann, 1976.

[3]R.W.Kates, Ed., "Managing Technological Hazard: Research Needs and Opportunities," Monograph 25, Institute of Behavioral Science, University of Colorado, Colo., 1977.

[4]A.B. Lovins, pp.104-115 of (6), below.

[5]T.B. Cochran, pp.93-94 of (6), below.

[6]S.M. Barrager, B.R. Judd and D.W. North, "The Economic and Social Costs of Coal and Nuclear Electric Generation," NSF 76-502, National Science Foundation, Washington, D.C., 1976

[7]I.R. Hoos, "Assessment of Methodologies for Radioactive Waste Manage-

ment," in W.P. Bishop, et.al., *Essays on Issues Relevant to the Regulation of Radioactive Waste Management*, NUREG-0412, U.S. Nuclear Regulatory Commission, Washington, D.C., 31-46, 1978.

[8]K.D. Tocher, "Planning Systems," *Philosophical Transactions of the Royal Society of London A.* 287, 425-441, 1977.

[9]J.W. Lathrop, ed., "Development of Radiological Performance Objectives, Interim Results: Trade-offs in Attitudes Toward Radioactive Waste,"UCID-17925, Lawrence Livermore Laboratory, Livermore, Ca., 1978.

[10]S.R. Watson and V.N. Campbell, "Radiological Performance Objectives for Radioactive Waste Derived from Public Values," Final Report PR78-10-80, Decisions and Designs, Inc., McLean, Va., 1978.

[11]F.R. Farmer, "Reactor Safety and Siting: A Proposed Risk Criterion," *Nuclear Safety 8*, 539-548, 1967.

[12]C. Starr, "Social Benefit Versus Technological Risk," *Science 165*, 1232-1238, 1969.

[13]E.M. Clark and A.J. Van Horn, "Risk-Benefit Analysis and Public Policy: A Bibliography", Brookhaven National Laboratory, Upton, N.Y., 1976.

[14]R. Papp, P.E. McGrath, L.D. Maxim, and F.X. Cook, Jr., "A New Concept in Risk Analysis for Nuclear Facilities," *Nuclear News, pp.* 62-65, November, 1974.

[15]R.A. Howard, "Social Decision Analysis," *Proceedings of the IEEE 63*, 359-371, 1975.

[16]W. Edwards, "How to Use Multi-Attribute Utility Measurement for Social Decision-Making," *IEEE Transactions on Systems, Man and*

Cybernetics SMC- 7, 326-339, 1977.

[17]D.A. Seaver, "Assessment of Group Preferences and Group Uncertainty for Decision Making," Technical Report 001597-4-T, Social Science Research Institute, University of Southern California, Los Angeles, Ca., 1976.

[18]R.L. Keeney and C.W. Kirkwood, "Group Decision Making Using Cardinal Social Welfare Functions," *Management Science 22*, 430-437, 1975.

[19]National Academy of Sciences Advisory Committee on the Biological Effects of Ionizing Radiation, "The Effects on Populations of Exposure to Low Levels of Ionizing Radiation" (known as the BEIR Report), National Research Council, Washington, D.C., 1972.

[20]International Commission on Radiological Protection, *Annals of the ICRP #26*, Pergamon Press, Oxford, 1977.

[21]B. Fischhoff, P. Slovic, S. Lichtenstein, S. Read and B. Combs, "How Safe is Safe Enough? A Psychometric Study of Attitudes Towards Technological Risks and Benefits," *Policy Sciences 8*, 127-152, 1978.

[22]R.E. Goodin, "Uncertainty as an Excuse for Cheating our Children: The Case of Nuclear Wastes," *Policy Sciences 10*, 25-43, 1978.

[23]B. Fischhoff, P. Slovic and S. Lichtenstein, "Knowing What you Want: Measuring Labile Values," in T. Wallsten, *Cognitive Processes in Choice and Decision Behavior*, Lawrence Erlbaum, Hillsdale, N.J., 1980.

[24]R.L. Keeney and H. Raiffa, *Decisions With Multiple Objectives: Preferences and Value Tradeoffs*, Wiley, New York, N.Y., 1976.

[25]U.S. Nuclear Regulatory Commission, *Reactor Safety Study: An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants*,

WASH-1400 (NUREG 75/014), U.S. Nuclear Regulatory Commission, Bethesda, Md., 1975.

[26]Ad Hoc Review Group, H.W. Lewis, Chairman, *Risk Assessment Review Group Report to the U.S. Nuclear Regulatory Commission*, (known as the Lewis Report), NUREG/CR-0400, U.S. Nuclear Regulatory Commission, Bethesda, Md., 1978.

[27]M. Kamins, "A Reliability Review of the Reactor Safety Study," Rand Paper P-5413, The Rand Corporation, Santa Monica, Ca., 1975.

[28]H.J. Otway, D. Maurer and K. Thomas, "Nuclear Power: The Question of Public Acceptance," *Futures*, pp. 109-118, 1978.