Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Using Values in Operations Research

Ralph L. Keeney,

To cite this article:
Ralph L. Keeney, (1994) Using Values in Operations Research. Operations Research 42(5):793-813. https://doi.org/10.1287/opre.42.5.793

OR FORUM

USING VALUES IN OPERATIONS RESEARCH

RALPH L. KEENEY

University of Southern California, Los Angeles, California

(Received June 1994; accepted June 1994)

Values pervade the field of operations research. Expressed as objectives, goals, criteria, performance measures, and/or objective functions, they are necessary in theoretical operations research models and in applications. Because of their critical role, it is useful to develop these expressions of values from basic principles. This paper outlines how to identify values for a specific decision problem, how to structure these values to facilitate thinking and analysis, and how to quantify values. Since values provide the basis for interest in a problem, these same values should guide all of our effort on that problem. Two important uses of values are to create better alternatives for decision problems and to define decision problems that are more appealing than those that confront us. On another level, the operations researcher's values are crucial in selecting the research and applications that he or she pursues. I illustrate this with a brief summary of a few projects concerning both life-threatening risks and the storage of nuclear waste. The presentation concludes with a challenge to operations researchers to consider devoting some effort and talent to what I think of as the mega-risks facing our country.

You cannot do operations research (OR) without values. Operations research is intended to improve decisionmaking; and values, indicating what one wants to achieve, are essential for guiding decisionmaking.

Morse and Kimball (1951) define operations research as "a scientific method of providing executive departments with a quantitative basis for decisions." Why should anyone bother with OR or, for that matter, bother to make decisions at all? The only answer is values. You want better, rather than worse, consequences and better and worse are based on values. Values are what we fundamentally care about in decisionmaking. Alternatives are simply means to achieve our values.

In an operations research model, values are used to specify the criteria for describing the consequences of alternatives. Values are necessary to build a quantitative objective function, which provides the basis for evaluating alternatives. Values are also used in deciding how to build any OR model. You want it to be good and useful and these desires indicate different values in different decision contexts. Perhaps most important, but sometimes not explicitly recognized, values—our values as operations researchers—are essential in choosing and defining the problems that we address.

This paper discusses many aspects of values in operations research, which has been the main focus of my work in the field. Values were critical in introducing me to OR and in all my professional decisions. After briefly describing some personal history, I summarize an approach to and some results for quantifying values into objective functions. Then I discuss several important concepts from value-focused thinking that concern how to systematically develop the objectives to be used in quantification, how to create better alternatives for any decision problems that you face, how to identify the decisions to work on, and how to use values to guide and integrate all these activities. Following this is an outline of four problems that I have chosen to work on in the recent past. I conclude with a challenge to you to think about what I believe are two of the most fundamental problem areas facing our nation.

PERSONAL HISTORY

It was Christmas vacation of 1963. I was a sophomore in a general engineering program at UCLA. My

Subject classifications: Decision analysis, multiple criteria: structuring and quantifying objectives. Government, regulations: examination of indirect consequences. Professional, OR/MS implementation: using values to select problems.
Area of review: OR FORUM.

0030-364X/94/4204-0793 $01.25

cousin Jim, an applied mathematician at Boeing, and the only one of my relatives who had gone to college, was visiting. I asked Jim what field I should specialize in. Jim replied that "people in operations research at Boeing are specifying their own hours and salaries." This was certainly consistent with my values. I looked into OR and took a few courses in it during my junior and senior years. Then I had the good fortune to be accepted into the graduate OR program at MIT.

Professor Philip Morse was the founder of the Operations Research Center at MIT and served as its Director through 1968. Professor Al Drake was Associate Director during my time as a student from 1966–1969. To choose a master's thesis topic, I spoke with each of the faculty affiliated with OR at MIT. One topic Al Drake mentioned was "how people should make decisions." That sounded interesting, so I began reading some literature about it and soon thereafter heard a lecture by Professor Howard Raiffa on the topic. I took Howard's course on decision analysis at Harvard and Al supervised my master's thesis on quantifying objective functions with two objectives.

During my second year at the Operations Research Center, I took a year-long doctoral seminar with Professors Raiffa and Richard Meyer at Harvard and began my dissertation on quantifying objective functions with more than two objectives. I was fortunate in being able to arrange with MIT that Howard serve as my dissertation advisor. I am even more fortunate in that I continue to work with Howard. He moved through the roles of teacher, advisor, mentor, colleague, and friend without losing any of them.

In addition to Howard and Al, I had the privilege of having Professor Morse on both my doctoral examination committee and my dissertation committee. After graduating and joining the MIT faculty in 1969, Phil—as he was now known to me—and I organized a one-week program on OR for public systems in the summer of 1970. In 1972, with Al, we three edited the book *Analysis of Public Systems* (Drake, Keeney and Morse 1972) that included the work presented in the program. I have fond memories and learned some important lessons from Professor Morse in these activities and numerous other meetings at the Operations Research Center.

Sometime during my first year at the Center, I realized that I had not really had much opportunity to talk with Professor Morse. One of my values was to get to know him. I made an appointment to talk with him about OR in general and a few specific topics, one of the latter being operations research faculty meetings. I suggested that it might be reasonable to have students attend most of these meetings. My proposal was based on values that were common to the faculty and students. Opening the meetings would improve communications between faculty and students, it would help students understand what faculty members do, and the faculty may get some work done by students, which would be good for both groups. For the few matters where students should not be present, we could simply be asked to leave. Professor Morse quietly considered this suggestion for about 15 seconds, and then said, "That's a good idea. How would you like to attend our next meeting?" From that meeting on, students were always welcome at the faculty meetings, apparently to everyone's benefit.

Two stories that Phil told about his work in operations research during World War II have influenced my choice of problems to study. One was that his team never worked on a problem unless they thought they could identify ways to yield at least a 300% improvement in operations. This said to me that the values of the operations researchers are important for choosing problems, and that the operations researcher should spend time to create alternative solutions to that problem. The other story was that his team wanted to think about the process of antisubmarine warfare before reading about it or analyzing any data on it. The insight here is that the operations researcher should be involved in identifying the problem areas to address. In this, your values are crucial.

MODELS OF VALUES

When facing a complex decision, most analysts argue that a model will lend insight into the complexity and will lead to better decisions. Furthermore, to gain the insight desired from such a model, it should focus on the basis for the complexity—often conflicting and intertwined values.

Most models of decision problems can be broken into two parts. The first part relates the available alternatives to the possible consequences of each alternative. These consequences describe the degree to which objectives are achieved. The second part evaluates the relative desirability of the consequences by means of a value model. One then combines information from the two parts of the overall model to derive the relative desirability of alternatives from the relative desirability of the consequences and the likelihoods that the various consequences will result from each alternative. The likelihoods of the consequences are quantified using data and judgments about facts. The desirabilities of the consequences are quantified using judgments of value. This distinction is important both because the techniques for addressing the two parts of the decision model are different and because the appropriate individuals for judging the two parts may be different.

Notation for a General Decision Model

It will facilitate our discussion if I introduce a minimum of notation. Let A_j, $j = 1, \ldots, J$, be a set of alternatives identified for a decision problem. Let O_i, $i = 1, \ldots, n$, be a set of objectives identified for that decision problem. The factual part of the decision model is to determine how well the alternatives perform in terms of the objectives. For this, we need a set of measures to indicate the degree to which the objectives are met. These measures are called attributes and are designated by X_i, $i = 1, \ldots, n$. We will use x_i to indicate a specific amount of X_i. Thus, for instance, if objective O_1 is "to maximize profits," the attribute X_1 might be "profits in millions of dollars," and a level x_1 might be 168, meaning 168 million dollars. With this notation, a possible consequence of an alternative can be written $(x_1, \ldots, x_n)$ or simply as x.

A value model, which is often simply referred to as an *objective function*, assigns a number $v(x)$ to each consequence x such that these numbers can be used to determine preferences for the alternatives. Different types of objective functions are appropriate for decision problems with particular characteristics. If there is no uncertainty relating the alternatives to the consequences and if one is interested only in ranking alternatives, then the value model need only assign higher numbers to preferred consequences. The resulting objective function is called an *ordinal value function*. Usually more information about the desirability of the alternatives is wanted. For example, one may be interested in the strength of preference for some alternatives over others. In this situation, the objective function should be a *measurable value function* constructed so that the differences in the values derived for the alternatives have a meaning beyond the mere ranking of preferences (see Dyer and Sarin 1979).

In many decision problems there are uncertainties about what consequences might result from any chosen alternative. In this case, the value model should be a *utility function*. This means that if one accepts a set of logical principles to guide the decision at hand (see, for instance, Pratt, Raiffa and Schlaifer 1964), the expected utility derived for each alternative indicates its relative desirability. An alternative with a higher expected utility should be preferred to one with a lower expected utility. However, the difference in expected utilities has no meaning unless the utility function also happens to be a measurable value function.

Before I proceed, it is worth mentioning one remarkable fact. Many complex decision problems have significant potential consequences, including, for example, costs of hundreds of millions of dollars, or fatalities, or large-scale environmental degradation. The only reason for taking an interest in such problems is that some consequences may be much better than others, and so some alternatives may be much better than others. Yet the amount of time usually taken to articulate appropriate values for a decision problem is minuscule relative to the time used to address other aspects of the problem. The objective function might be chosen in an hour with very little thought, and yet several person-years of effort and millions of dollars may be used to model the relationships between alternatives and consequences and to gather information about those relationships. Since the sole reason for caring about the problem is based on values, it would not seem unreasonable to use a fraction of those resources to understand and appropriately structure and quantify the relevant values in building a value model.

General Procedure to Build a Value Model

A value model, naturally enough, is a model of values with quantitative relationships. The procedure to build this model is essentially the same as for any model—as is the motivation, namely to lend some insight about a complex situation to complement intuitive thinking.

The key to building a value model is the set of objectives and the corresponding attributes. One must first identify a set of objectives appropriate for the decision problem. Then, one needs to define attributes for measuring the degree to which these objectives are met. For example, two of many objectives for evaluating sites for a hazardous waste dump might be to minimize cost and to minimize environmental damage. For the sites under consideration, corresponding attributes may be selected as millions of dollars and acres of forest destroyed.

A general structure must be developed next to combine the various attributes in some proper manner to reflect preferences. This is analogous to building a model relating profits to sales, unit selling price, and fixed and variable costs of producing a product. To combine the various attributes, one uses independence concepts analogous to the concept of probabilistic independence in building factual models. To structure relationships among different levels of an attribute, concepts of strength of preference and risk attitude are relevant. For instance, whether an individual is risk averse or risk neutral is reflected in the functional form of the value model.

The general structure of the value model has parameters that can be determined by value judgments. These value judgments are the "data" that one needs to construct the model. The data, which rest in the minds of decisionmakers or individuals knowledgeable about a given problem, is gathered by eliciting judgments from those people. This situation is analogous to many scientific problems where the necessary knowledge to parameterize a model is "out there" and needs to be collected. If the information is about geology, one digs holes in the ground to collect data. If the information is about values, one "digs holes" into the mind to collect data.

Finally, once the first-cut version of the model is complete, it is necessary to check its reasonableness by examining how it performs in situations that are well understood. Then, modifications may be necessary until one believes that the model adequately reflects the appropriate values.

Specific Procedures to Build a Value Model

My original research, which included both my master's thesis and doctoral dissertation, concerned the quantification of values. The following is a summary of much of this work.

> An optimum decision in a problem involving more than one objective or goal is one that maximizes some function of the effectiveness of the decision relative to the objectives and the relative importance of the objectives. In effect, an optimum decision is one which maximizes 'weighted effectiveness' where the weighting factors are the values of the objectives. A practical procedure for obtaining from decisionmakers a measure of the relative importance of their objectives is described and illustrated. The technique is applicable to problems involving any number of objectives and any number of decision makers.

I did not write this summary. It is the abstract for a paper titled, "An Approximate Measure of Values," by Churchman and Ackoff presented at the First Annual Meeting of the Operations Research Society of America held in May 1953 in Cleveland. My choice of words to describe my work would be slightly different, but the problem that I addressed is the same.

There are now many procedures and much technical information for building value models. They all rely on using independence concepts to combine the achievement on different objectives into a value model. The books by Keeney and Raiffa (1993) and von Winterfeldt and Edwards (1986) have more details on these matters than most readers would care to know.

Intuitive definitions of four important independence concepts are:

Additive Independence. Attributes are additive independent if preferences for the consequences depend only on the individual levels of the separate attributes and not on the manner in which the levels of the different attributes are combined.

Preferential Independence. A pair of attributes is preferentially independent of other attributes if preferences for levels of these two attributes do not depend on the levels of any of the other attributes (i.e., the indifference curves over the pair of attributes do not depend on levels of the other attributes).

Utility Independence. An attribute is utility independent of other attributes if preferences for risky situations (i.e., lotteries) involving probabilities of different levels of that attribute do not depend on a fixed level of any of the other attributes.

Weak-Difference Independence. An attribute is weak-difference independent of other attributes if preference differences between pairs of levels of that attribute do not depend on a fixed level of any of the other attributes.

Using these independence concepts, one can derive structures for the overall value function. Two of the most common structures are the additive and multiplicative value models.

Additive Value Model. If all combinations of attributes are additive independent, then the utility function u must have the additive form

$$u(x_1, \ldots, x_n) = \sum_{i=1}^{n} k_i u_i(x_i), \qquad (1)$$

where u and the u_i are utility functions scaled from 0 to 1 and the k_i are positive scaling constants summing to 1.

The original proof and details of this result are found in Fishburn (1965).

Multiplicative Value Model. If each pair of attributes is preferentially independent of the others and if one attribute is utility independent of the others, then the utility function u must have the multiplicative form

$$1 + ku(x_1, \ldots, x_n) = \prod_{i=1}^{n} [1 + kk_i u_i(x_i)] \qquad (2)$$

or the additive form where u and the u_i are utility functions scaled from 0 to 1, the k_i are positive scaling constants, and k is an additional scaling constant.

The original proof of this result for three or more attributes is found in Keeney (1974). With two attributes, the value model (2) is appropriate if each attribute is utility independent of the other (Keeney 1968). If weak-difference independence is substituted for utility independence in the result above, the resulting value model will be either a multiplicative or additive measurable value function (see Dyer and Sarin 1979).

To examine the appropriateness of any independence condition, one considers specific cases with the individual whose preferences are being assessed to see whether they violate the independence condition. For instance, to examine whether the attribute pair X_1 and X_2 is preferentially independent of the other attributes, one first identifies pairs of attribute levels of X_1 and X_2, that is (x_1, x_2), that are indifferent to each other. If this indifference is maintained regardless of the levels of the other attributes, X_3 through X_n, then the preferential independence condition holds.

It is worth noting that whenever a specific independence condition is not appropriate, an objective in addition to those articulated for the problem is relevant. For instance, suppose for a decision problem with two objectives that additive independence did not hold, meaning that the decision maker did care how the levels of the different attributes were combined. This implies that the overall value could not be determined by simply adding up the value due to the two objectives, appropriately weighted, as indicated by (1). However, suppose that each attribute was utility independent of the other, implying the specific case of value model (2) that can be written

$$u(x_1, x_2) = k_1 u_1(x_1) + k_2 u_2(x_2) + k_3 u_1(x_1) u_2(x_2). \qquad (3)$$

Value model (3) is clearly additive; three components are added to get the overall utility. The first two terms represent the contributions from objectives O_1 and O_2; the third term captures a third, 'hidden,' objective. Identifying what value it is meant to capture can provide an important insight for examining the decision problem.

The Most Common Critical Mistake

One mistake is very commonly made in constructing value models. Unfortunately, that mistake is sometimes grounds for poor decision making. It is illustrated in the context of an air pollution problem where one is concerned about the costs of regulating air pollution emissions and about the concentration of pollutants. Administrators, regulators, and members of the public are asked questions such as the following. Which is more important, costs or pollutant concentrations? And how much more important is the more important attribute? A reply might be that pollutant concentrations are three times as important as costs. While this statement may make sense as a sentiment, it is completely useless for building a model of values. Would the statement mean that lowering air pollutant concentrations by one part per billion would be worth a cost of \$2 billion in a metropolitan area? The answer is probably, "Of course not." Indeed, that answer would come from someone who had just stated that pollutant concentrations were three times as important as costs. When asked to clarify the apparent discrepancy, that person would naturally state that the decrease in air pollution was very small, only one part in a billion, and the cost was a very large \$2 billion. The point should be clear. One must know by how much air pollution concentration will change and how much the change will cost to logically discuss and quantify the concept of relative importance.

This error causes obvious problems. It fails to provide the in-depth appraisal of values that important decision problems require. If we are concerned about the effect of pollutant concentrations on the health of the public and the expenditures of billions of dollars, I personally do not want some administrator to give two minutes of thought to the matter and state that pollutant concentrations are three times as important as costs. Such judgments are often made by the public, concerned groups, or legislators, and are then used inappropriately. Indeed, sometimes legislators pass laws such as the Clean Air Act in the United States, which essentially states that the health of the public is of paramount importance and that costs in achieving desired air pollutant levels should not be considered in setting standards. This is not practical or desirable in the real world. We could spend hundreds of billions of dollars and still be able to improve air quality with additional expenditures. This would be the case even if we could improve the "national health" by reducing the number of annual asthma attacks in the country by only five. In short, if the value tradeoffs are done properly and specify how much of one attribute is worth how much of another, the insights from the analysis are greatly increased and the likelihood of misuse of judgments is greatly decreased.

MAKING VALUES EXPLICIT

Strategic thinkers have long recognized the need to clarify values as a key step in developing strategic approaches (e.g., Franklin 1772). Peters and Waterman (1982) refer to their "one all-purpose bit of advice for management" in the pursuit of excellence as "figure out your value system." Values are clarified with an explicit statement of specific objectives. However, identifying and structuring objectives is a difficult task: ends are often confused with means, and objectives with targets or constraints or even alternatives; the relationships among objectives are unclear; and the concept of priorities within objectives is easily misconstrued. Clear objectives are very useful, but how should they be developed? The process requires considerable creativity in discussions with decisionmakers and those concerned with the decision. Simply listing objectives is shallow. There is a need for greater depth, clear structure, and a sound conceptual base in developing objectives for strategic decisions.

An objective is something that one wants to strive toward. It has three features: a decision context, an object, and a direction of preference. For example, one objective of a forest products company is to "minimize environmental impacts." The decision context here is harvesting natural resources, the object is environmental impact, and the preference is for less impact rather than more.

We distinguish between fundamental objectives and means objectives. Fundamental objectives concern the ends that decision makers value in a specific decision context; means objectives are methods to achieve ends. Fundamental objectives for strategic decisions—the broadest class of decisions facing an organization—are defined as strategic objectives. The strategic objectives provide common guidance for all decisions in an organization and form the basis for more detailed fundamental objectives appropriate for specific decisions.

Identifying Objectives

The most obvious way to identify objectives is to discuss the decision situation with decisionmakers and stakeholders. In this discussion, you respectfully interrogate the decision makers and stakeholders about their values pertaining to the decision situation. The techniques listed in Table I help stimulate the identification of possible objectives. There is redundancy in these techniques, and intentionally so: It is much easier to recognize redundant objectives when they are listed than it is to identify missing objectives.

When asking someone to express objectives, make it clear that what is first needed is simply a list of objectives without ranking or priorities. You might begin a discussion by asking, "What would you like to achieve in this situation?" To expand the list, ask, "If you had no limitations at all, what would your objectives be?" Similarly, ask about the elements of the desirable end result for the decisionmaker.

Many words trigger questions that help make implicit objectives explicit—words such as tradeoffs, consequences, impacts, concerns, fair and balance. For example, if a decisionmaker says "Tradeoffs are

Table I
Techniques to Use in Identifying Objectives

1. **A Wish List.** What do you want? What do you value? What should you want?
2. **Alternatives.** What is a perfect alternative, a terrible alternative, some reasonable alternative? What is good or bad about each?
3. **Problems and Shortcomings.** What is wrong or right with your organization? What needs fixing?
4. **Consequences.** What has occurred that was good or bad? What might occur that you care about?
5. **Goals, Constraints and Guidelines.** What are your aspirations? What limitations are placed upon you?
6. **Different Perspectives.** What would your competition or your constituency be concerned about? At some time in the future, what would concern you?
7. **Strategic Objectives.** What are your ultimate objectives? What are your values that are absolutely fundamental?
8. **Generic Objectives.** What objectives do you have for your customers, your employees, your shareholders, yourself? What environmental, social, economic, or health and safety objectives are important?
9. **Structuring Objectives.** Follow means-ends relationships: why is that objective important, how can you achieve it? Use specification: what do you mean by this objective?
10. **Quantifying Objectives.** How would you measure achievement of this objective? Why is objective A three times as important as objective B?

necessary," ask tradeoffs between what and what. If a decisionmaker says "The consequences should be fair," ask fair to whom, and what is fair. If the decisionmaker seems to stop and think, ask what those thoughts are. Responses to these questions may suggest potential objectives and provide a basis for further probing.

Hard thinking about a decision situation often begins only after at least two alternatives become apparent. Once one articulates the features that distinguish alternatives, one can identify some objectives. For example, if a site is to be chosen for an airport, one feature that differentiates alternatives might be disruption due to noise. This suggests the obvious objective of minimizing disruption from noise. You might also ask respondents to list particularly desirable and undesirable features of alternatives and use these to stimulate thought about objectives.

Structuring Objectives

The initial list of possible objectives will contain many items that indicate values but that are not really objectives. These will include alternatives, constraints, and criteria to evaluate alternatives. With some thought, each item on the list can be converted into objectives. Now the list of proposed objectives will include both fundamental objectives and means objectives. It is important to separate these types of objectives and establish their relationships by examining the reasons for each item on the list. Two important concepts are repeatedly used. One involves linking objectives through means-ends relationships; the other involves specifying fundamental objectives.

Repeatedly tracing ends objectives from specific means objectives should lead to at least one fundamental objective. For each objective, ask, "Why is this objective important in the decision context?" Two answers seem possible. One answer is that the objective is one of the essential reasons for interest in the situation: Then it is a fundamental objective. The other answer is that the objective is important because of its implications for some other objective. In this case, it is a means objective, and the response to the question identifies another objective. The "why is it important?" test must be given here in turn to ascertain whether this is a means objective or a fundamental objective.

Consider a decision involving the transportation of nuclear waste. One objective may be to minimize the distance the material is transported by trucks. The question should be asked, "Why is this objective important?" The answer may be that shorter distances would reduce both the chances of accidents and the costs of transportation. However, it may turn out that shorter transportation routes go through major cities, exposing more people to the nuclear waste, and this may be deemed undesirable. Again, for each objective concerning traffic accidents, costs, and exposure, the question should be asked, "Why is that important?" For accidents, the response may be that with fewer accidents there would be fewer highway fatalities and less accidental exposure of the public to the nuclear waste. And the answer to why it is important to minimize exposure may be to minimize the health impacts due to nuclear waste. To the question "Why is it important to minimize health impacts?", the response may be that it is simply important. This indicates that the objective concerning impacts on public health is a fundamental objective in the decision context. Figure 1a presents part of a means-ends objectives network that elaborates several means by which health and safety impacts occur.

Another important concept in identifying fundamental objectives is that of specification. The intent is to break an objective into its logical parts. Consider

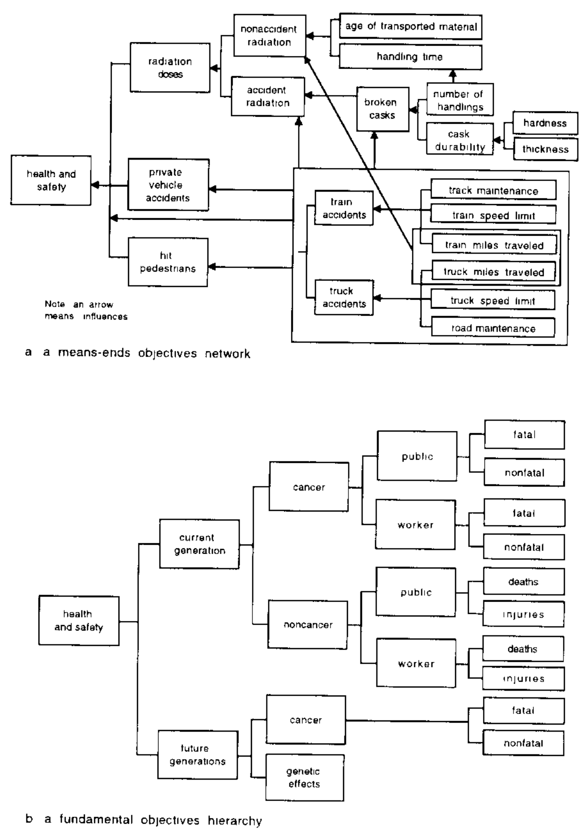

Figure 1. Objective structures for the safety impacts of transporting nuclear waste.

Published in Operations Research on October 01, 1994 as DOI: 10.1287/opre.42.5.793.
This article has not been copyedited or formatted. The final version may differ from this version.

the fundamental objective pertaining to health and safety. To make this more specific, ask, "What health and safety impacts should be minimized?" The response, which may be one impact or many, should lead to clarification of the objective and a better focus for thinking and action.

A portion of the fundamental objectives hierarchy pertaining to health and safety is illustrated in Figure 1b. The first level of specifying the impacts on health and safety of transporting spent nuclear fuel is impacts on the current generation and impacts on future generations. The impacts on future generations are further specified as genetic effects and cases of cancer, some that will be fatal and some that will not. Impacts on the current generation are divided into cancer cases and other (noncancer) effects, and are categorized by whether they affect the public or workers in the nuclear industry. The cancer implications are further specified as fatal or not, and the noncancer implications are specified as resulting in deaths or in serious injuries.

It is important to recognize that different types of judgments are necessary to structure fundamental objectives hierarchies and means-ends objectives networks. Value judgments are required to construct fundamental objectives hierarchies, and judgments about facts are required to construct means-ends networks. Quite simply, deciding what is important requires value judgments. Deciding how to achieve a higher-level objective requires factual knowledge. This difference means that for some decision situations it may be appropriate to have different individuals build the two types of objectives structures. For example, consider problems of public policy where solutions involve technical complexity. The public's values, or values expressed for the public by representatives (such as legislators or regulators), are appropriate in the construction of the fundamental objectives hierarchy. But when it comes to constructing the means-ends objectives network, individuals with expertise about technical or factual aspects of the decision situation are often better qualified than the public or its representatives.

VALUE-FOCUSED THINKING

Values are fundamental to all that we do; thus, values should be the driving force for our decision making. They should provide the basis for the time and effort we spend thinking about decisions. But they do not. Instead, decision making usually focuses on alternatives. Decision problems are thrust upon us by the actions of others—competitors, customers, government, and stakeholders—or by circumstance: recessions and natural disasters. Faced with a decision problem, the so-called solving begins. Typically, the decisionmaker defines the problem as a choice among identified alternatives and only afterwards considers objectives or criteria to evaluate them. I refer to this standard problem-solving approach as alternative- focused thinking.

Focusing on alternatives is a limited way to think through decision situations. It is reactive, not proactive. If you wish to be the master of your decision making, it makes sense to have more control over the decision situations you face. You do not control decision situations that you approach through alternative-focused thinking. This standard mode of thinking is backwards, because it puts the cart of identifying alternatives before the horse of articulating values.

It is values that are fundamentally important in any decision situation. Alternatives are relevant only because they are means to achieve values. Thus, although it is useful to iterate between articulating values and creating alternatives, the principle should be "values first." This manner of thinking, which I refer to as value-focused thinking, is a way to channel a critical resource—hard thinking—to lead to better decisions (Keeney 1992).

The central role of thinking about values is illustrated in Figure 2. Value-focused thinking helps

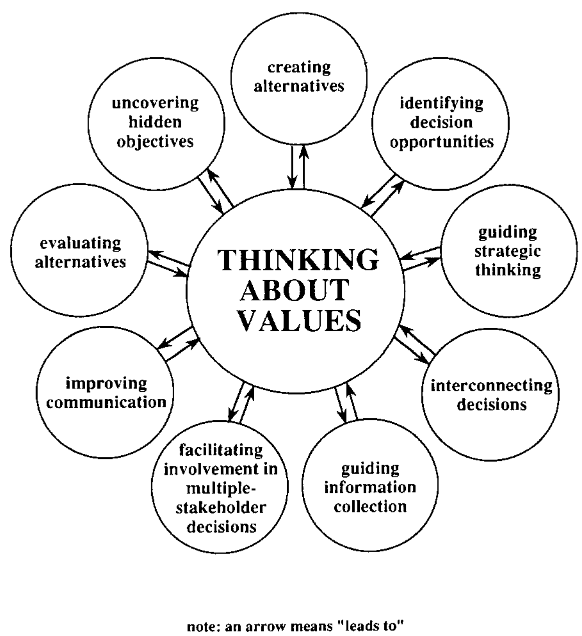

Figure 2. The crucial role of thinking about values in value-focused thinking.

uncover hidden objectives and leads to a more productive collection of information. It can improve communication among parties concerned about a decision, facilitate involvement of multiple stakeholders, and help to coordinate interconnected decisions. This all aids the evaluation of alternatives. But the greatest benefits of value-focused thinking result from the guidance provided for creating better alternatives for any decision problem that you face, and for identifying decision situations more appealing than the decision problems that confront you. These better decision situations that you create for yourself should be thought of as decision opportunities, rather than as decision problems.

Value-focused thinking has been applied recently in several important decisions. These include the operations of the British Columbia Hydro and Power Authority, a publicly-owned electric utility with billions of dollars in annual sales that supplies power to over 90% of British Columbia (Keeney and McDaniels 1992); the internal review of the policies of Conflict Management, Inc., a consulting organization concerned with negotiations and dispute resolution, located in Cambridge, Massachusetts (Keeney 1994); and the possible development of coal reserves discovered in a remote rainforest of Sabah, a State of Malaysia located on the island of Borneo (Gregory and Keeney 1994).

Creating Alternatives

The range of alternatives people identify for a given decision situation is often unnecessarily narrow. There are several reasons for this. There is a tendency in all problem solving to move quickly away from the ill-defined to the well-defined, from constraint-free thinking to constrained thinking. There is a need to feel, and perhaps even measure, progress toward reaching a "solution" to a decision problem. To get that feeling of progress, one often quickly identifies some viable alternatives and proceeds to evaluate them, without making the effort to broaden the search for alternatives.

The first alternatives that come to mind in a given situation are the obvious ones, those that have been used before in similar situations and those that are readily available. Once a few alternatives—or perhaps only one, such as the status quo—are on the table, they anchor thinking about others. Assumptions implicit in the identified alternatives are accepted, and new alternatives, if any are generated, tend to be merely variants of those already identified. Truly creative or different alternatives remain hidden elsewhere in the mind, unreachable by mere tweaking. Only deep and persistent thought can jar them into consciousness. Focusing on the values that should be guiding the decision situation frees one's thinking from narrowly defined alternatives and makes the search for others a creative and productive exercise.

Numerous guidelines aid the search for alternatives, or more precisely, the search for good alternatives. Alternatives should be created that best achieve the values specified for the decision situation. Both the qualitative objectives and any quantitative statements of values (e.g., priorities) should be systematically probed to initiate creative thought. Ideally, you want to create the best possible alternative using the least amount of time, effort, and resources. But realistically, in complex decision situations, the practical aim in creating alternatives is to generate a set of very promising ones.

The objectives in the fundamental objectives hierarchy list all aspects of consequences that are important in the context of a decision. Hence, thinking about how to better achieve these objectives can suggest alternatives. To begin, it is useful to focus on one objective at a time and think of alternatives that might be desirable if it were the only one. You should consider every objective, regardless of its level in the hierarchy. This exercise is likely to produce a broad range of alternatives, most of which would rank rather poorly on objectives other than the one for which they were devised. If this is not the case, you have not been very creative in generating the alternatives.

The next step is to consider objectives two at a time and try to generate alternatives that contribute to both. These are likely to be either refinements or combinations of alternatives created using single objectives. Then take three objectives at a time, and so on, until all objectives are considered together. Next, examine the alternatives you have generated to see whether any of them can be combined into a single alternative.

The means objectives also stimulate thinking about alternatives. Essentially any alternative that influences one of the means objectives should influence at least some of the associated fundamental objectives. More complicated decision situations will tend to have more means objectives, which should suggest more alternatives. Consider the rather simple fundamental objective of a firm to maximize profits from the sale of a product. Means objectives, in this case, might concern sales volume, price of the product, cost of manufacturing the product, cost of distributing the product, and overhead cost. Alternatives could be created that influence each of these.

A decisionmaker's broadest objectives are the strategic objectives. These provide the foundation for creating any alternatives or identifying any decision opportunities based on values. It can be useful to ponder the wide-open question of what can be done to contribute to the achievement of each strategic objective. Applying the ideas outlined to each strategic objective should suggest some worthwhile alternatives. Using the strategic objectives in this way, without a decision context, is closely related to identifying decision opportunities.

Descriptive research on decision making has made it clear that the processes typically used to arrive at decisions are often not as logical or systematic as we would like them to be (see Russo and Schoemaker 1989). Thus, at the end of a decision process, when you are about to choose an alternative, think about new alternatives once more. At this stage, you know how good that about-to-be-chosen alternative is. You should also have a good idea of the time, money, and effort necessary to implement it. Now carefully search for a new alternative that can better achieve the fundamental objectives by using the same or fewer resources in a different way.

Decision Opportunities

Who should be making your decisions? The answer is obvious: You should. Well then, who should be deciding what decision situations you face and when you face them? The answer here is the same: You should. At least you should control far more of the decision situations you face than many of us do. That control may influence the achievement of your objectives far more than does controlling the alternatives selected for the decisions.

Decisionmakers usually think of decision situations as problems to be solved, not as opportunities to be taken advantage of. Thus, it is not surprising that decisionmakers do not systematically hunt for decision situations. Who needs yet another problem? One of the main precepts of value-focused thinking is that, perhaps, you need another "problem"; that a decision problem may not be a problem at all but an opportunity. A decision opportunity may alleviate some other decision problems or, perhaps, allow you to avoid future problems. In this sense, recognizing and following up on decision opportunities is analogous to prevention, whereas dealing with decision problems is analogous to cure.

There are two ways to create decision opportunities. One is to convert an existing decision problem into a decision opportunity. Often this involves broadening the context of the problem (see Nadler and Hibino 1990). The other way to create decision opportunities is from scratch. You use your creative genius, which can be stimulated by value-focused thinking, to examine whether and how you can better achieve your objectives.

The strategic objectives of an organization are the foundation from which decision opportunities can be identified. But how many organizations have written down and organized their strategic objectives? In how many organizations do the employees know and understand the organization's strategic objectives? The answer to both questions is very few. Many organizations have mission statements and, perhaps, lists of organizational goals, but these are often too vague to guide the creation of decision opportunities. The choice to state and clarify strategic objectives is itself a decision opportunity.

The decisionmaker should develop procedures to initiate the search through the strategic objectives for decision opportunities. Some of these procedures should be independent of decision situations currently being addressed. For example, it may be useful to set up a procedure of seeking potential decision opportunities on the first of every month. To put this procedure into operation, the decisionmaker must unambiguously state and structure its strategic objectives. Then, each month the decisionmaker should verify that the strategic objectives are appropriate and make any necessary changes. There should usually be no changes because we are talking about broad, high-level objectives. Then the decisionmaker should use each strategic objective to help stimulate thinking of decision opportunities to do things better.

Decision opportunities can be very helpful to you when you have no direct control over a decision that you care about. In an important class of such decisions, one stakeholder wishes to see a certain alternative selected, but another stakeholder has the power to make the decision. A firm wants its proposal to supply a product to another firm accepted, a government wants another government to sign an agreement, an executive wants his application for a sabbatical granted by the organization, and so on.

Suppose that you are the stakeholder who wants a particular alternative selected by another stakeholder, whom I will now refer to as the decisionmaker. You should recognize your decision opportunity to take control of the decision situation. Rather than simply allow the decisionmaker to choose an alternative that may not be the one you desire, you should create alternatives that modify your desired alternative so that it maintains its essential features for you and is better than the existing alternatives available to the decisionmaker.

The key to "solving" a problem of this nature is to view it from the perspective of the decisionmaker. First structure her values as much as possible. Indeed, she should have little objection to your efforts in this regard, as she may end up better off as a result. Then identify the undesirable impacts to her of your desired alternative relative to the status quo. These impacts probably affect the means objectives. Follow their implications through a means-ends objectives hierarchy to the fundamental objectives of the decisionmaker. Now you should be able to create modified alternatives that can improve matters in terms of her fundamental objectives while maintaining the key consequences desired by you. In many cases, the decisionmaker may prefer the modified alternatives to the status quo. It is worth noting that the strategic objectives of a decisionmaker are likely much broader than the set of fundamental objectives influenced by your desired alternative. Thus, an add-on alternative distinct from the original one may be used to "satisfy" her when used in conjunction with the original alternative.

From your perspective, the decision alternatives include the status quo and any others you can create. From the decisionmaker's perspective, the alternatives are the status quo and the alternative that you are suggesting. The problem may be characterized as an empathetic negotiation. You negotiate for the other side and make sure the decisionmaker gets enough to be willing to support your desired alternative, and you negotiate for your side as well. This situation may also be viewed as removing constraints to action. The interplay between the descriptive and prescriptive aspects of decision making in empathetic negotiations is interesting. You are essentially using your description of the decisionmaker's values both to create new alternatives and as a basis for negotiating her position prescriptively. Also, you are trying to balance overall impacts to the decisionmaker and overall impacts to you in a manner that is prescriptively reasonable to you, and that the decisionmaker will view as descriptively fair and responsible. In simple terms, your goal is to create a win-win alternative.

USING YOUR VALUES TO GUIDE YOUR WORK

Values should guide the problems that operations researchers work on and how they address those problems. The values must be those of the operations researcher. Sometimes they will be only means to achieve the objectives of one's employer; one might work to do the best job on a problem assigned by the boss. But usually, even in such cases, there is considerable flexibility in defining and pursuing the problem posed. For many operations researchers, especially those in universities, there are numerous decisions to be made about one's work. The fundamental role of values in guiding the professional path you follow is succinctly illustrated by a line from Lily Tomlin: "I always wanted to be somebody, but I should have been more specific."

Clarifying your values for work is the way to get specific about your professional life. It specifies what you hope to accomplish. Let me illustrate this simple but very important notion with some personal examples as a way of setting up my challenge to operations researchers.

My fundamental professional objectives, discussed in detail in Keeney (1992), are to maximize my enjoyment and my learning, provide service, enhance my career, and build good professional relations. Means objectives concern the selection of problems that are interesting, not overstudied by others, important (i.e., having potentially significant consequences), relevant to other stakeholders, relevant to many specific situations, and that have no clear path for the work. I want the chance of a breakthrough, a chance to frame the thinking of others, and a chance to influence action. Let me summarize four problems on which I have chosen to work in pursuit of my professional objectives.

Indirect Health and Safety Risks of Regulations

In the early 1980s, in response to a legal suit by the Sierra Club, the Environmental Protection Agency was considering whether it should regulate radionuclide emissions from coal-fired power plants. After its review of the relevant information, the EPA (1983) stated that the capital cost for the equipment needed would be $13 billion and the annual operational cost $3.4 billion. It also stated that this expenditure might have a net health benefit of one less fatal cancer per year of the one to two predicted cancers per year that might be associated with current emissions. This situation seemed an ideal candidate to make the point that proposed health and safety regulations should be evaluated in terms of their net health and safety benefits. Regulations themselves induce indirect health and safety impacts that need to be considered. In some cases, these can be more severe than those one wishes to prevent.

My colleague Detlof von Winterfeldt and I built a model to estimate the indirect health and safety risks associated with environmental regulations of power plants (Keeney and von Winterfeldt 1986). It began with an engineering and technological assessment of the requirements for meeting a regulation. This led to specific construction, operations, and financial activities necessary to comply with the regulation. These activities have indirect health and safety effects from construction, generation of replacement power, employment changes, and income changes.

Construction is a relatively risky occupation compared to other occupations or the absence of employment. Requiring additional construction activities adds workers to the construction force, thereby subjecting them to greater safety risks than they were subjected to previously. Because electricity must be provided to operate most pollution control equipment and because operating power plants lose efficiency as pollution is controlled, additional generating capacity is required, which will increase health risks.

To pay for the costs of complying with environmental regulations, electricity rates would likely rise. The utility industry is borrowing large sums of money to purchase and construct pollution control equipment, and replacement power facilities would also tend to drive interest rates up. Increases in interest rates and higher electricity prices increase costs to American business; and as a result, some jobs may be eliminated. On the other hand, jobs would increase in the regulatory industries and in the power industry as generating capacity is added. Evidence suggests that stress is induced by losing a job or changing a job, and that this causes health risks; but these might be reduced if more jobs are created than are lost. Therefore, unemployment/employment effects could either increase or decrease health risks.

The cost of environmental regulations and additional generating capacity will eventually be passed on to individual households. The mechanisms may include lower wages (or even no wages), higher interest rates, higher product costs, and higher electricity prices in real terms. Thus, there will be less real income available to consumers for health care, nutrition, and discretionary safety (for example, purchasing a home fire alarm). A positive relationship exists between higher real income and lower adverse health effects. Accordingly, this pathway is likely to lead to higher health risks.

This indirect risk model was applied to the radionuclide emissions problem. There was reasonable data for the first two pathways. For the additional construction required, we used fatality rates for the types of work that would be performed and calculated that 37 such fatalities could be expected. For replacement power, estimates for capacity lost through the loss of efficiency from the pollution equipment ranged from 0.9 to 2.2%. Applying this to U.S. coal power plants suggested the need for an additional 1,700 to 4,200 megawatts of capacity. Three different sources estimated that the annual fatalities from mining, processing, and transporting coal were in the range of 1.05 to 11.9 for each 1,000 megawatt-years of electricity. Multiplying the megawatts annually needed by the fatality rates suggests that 1.8 to 50 fatalities might occur annually. Any fatalities due to pollution would increase these figures.

These calculations were provided as testimony for the regulatory process concerning the proposed national radionuclide emissions standard. In addition, the logic and some evidence for the implications of unemployment and household income on health and safety were presented, although no estimates of fatalities were given because there was no reasonable model to calculate those in 1983. This situation is partially remedied as discussed below. The proposed regulation to limit radionuclide emissions from coal-fired power plants was not passed. Perhaps this work had some influence on this outcome.

Mortality Risks Induced by the Costs of Regulations

Regulations intended to reduce health and safety risks and thereby reduce premature deaths to Americans cost billions of dollars annually. Are they worth it? More precisely, how can we distinguish those proposed regulations that are worthwhile from those that are not? Part of the answer depends on whether the costs of a specific regulation induce more fatalities than the regulation itself is expected to avoid.

A wealth of evidence indicates that lower incomes are associated with higher mortality risks. Numerous authors, summarized in Lutter and Morrall (1994), show that the mortality rate for individuals with higher incomes is less than that for individuals with lower incomes. Reasons for this relationship relate to, among other things, better nutrition, better sanitation, better health care, better education, and better socioeconomic status—all of them easier to come by with money. This raises a key issue about whether the cost of a proposed regulation, which de facto reduces disposable income, would increase mortality risks and therefore produce more premature deaths than those purported to be saved by the proposed regulation.

Over the past 15 years, Wildavsky (1979, 1980, 1988) has clearly stated in qualitative terms the concept of the "richer is safer" argument. That argument, in its simplest form, is that the costs of regulations are borne by individuals. Intermediaries such as the government or companies have no options that do not pass costs on to individuals. Thus, at least temporarily, individuals are poorer in the sense that they have less disposable income for other purposes. These individuals use, on average, additional disposable income in ways that reduce their health and safety risks and therefore reduce fatalities. Thus, regulatory costs induce fatalities.

A model to quantify the richer is safer argument and estimate cost-induced fatalities is developed in detail in Keeney (1990). It considers the situation where the total cost of a regulation is felt in a single year and where the monetary costs of that regulation are redistributed exactly as they were charged to the individuals and society (i.e., there is no change in the distribution of income). The model has three fundamental components that address the following:

$r(x)$ = *the annual probability of death for an individual with income x*;

$f(x)$ = *the probability distribution for the annual income of individuals prior to paying for the costs of the regulation*; *and*

$c(x)$ = *the relative one-time cost of the regulation to an individual with income x*.

The first component was developed separately using data from two main sources, Kitagawa and Hauser (1973) and Frerichs et al. (1984). Both suggest that the curve in Figure 3 can be used to relate annual mortality risk and income. This exponentially decreasing curve has three intuitively appealing properties: the mortality risk increases as income decreases; the increases in mortality risk due to a specific income reduction are greater at lower income levels; and the effect of income on mortality risk is relatively insignificant at high income levels.

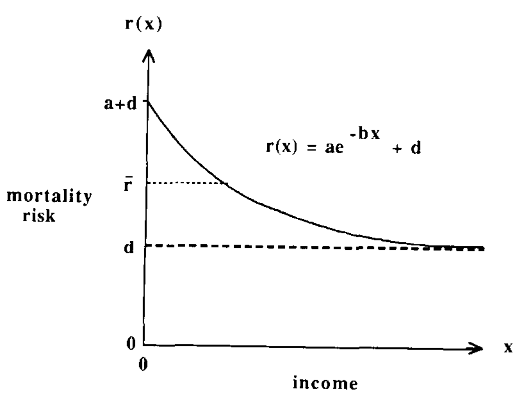

Figure 3. An aggregate model relating annual income to annual mortality risk.

A function representation for $r(x)$ with these properties is

$$r(x) = ae^{-bx} + d, \tag{4}$$

where a, b, and d are positive constants and x represents income. If x is very large, then $r(x) = d$, so d can be thought of as the mortality risk of the wealthy, or the residual mortality risk that is not reduced with additional income. When $x = 0$, then $r(x) = a + d$, which is the highest mortality risk. The constant a is the amount of mortality risk that can be influenced by income. The constant b, which has no simple interpretation, concerns the range of income over which there is an influence on mortality. The average mortality risk, labeled $\bar{r}$ in Figure 3, lies between mortality risks d and $(a + d)$.

Information to develop the second component of the basic model concerning the annual income of individuals is available from the Bureau of the Census (1981). A probability distribution that can be used to closely fit annual income data is the gamma distribution

$$f(x) = \frac{\lambda^{\alpha} x^{\alpha - 1} e^{-\lambda x}}{(\alpha - 1)!} \tag{5}$$

with positive parameters α and λ, where the factorial designates the gamma function when α is not an integer. The mean $\bar{x}$ (i.e., average income) and standard deviation σ_x of (5) are

$$\bar{x} = \alpha/\lambda \tag{6}$$

and

$$\sigma_x = (\alpha/\lambda^2)^{1/2}, \tag{7}$$

respectively, which can be calculated from the actual income data and substituted into (6) and (7) to solve for the income distribution parameters.

The gamma distribution has four intuitively appealing properties for an income distribution: the range is from 0 income on up; it has a single mode; the distribution is skewed to the left, implying that more than half the individuals have less than average income and that some individuals have several times the average income; and it can reasonably approximate a large number of possible income distributions by varying its two parameters α and λ.

For different regulations, regulatory costs may be borne very differentially by individuals in society. The manner in which costs are allocated among those individuals depends on the complex workings of our economy and the intermediaries (e.g., firms, government agencies). In the basic model, three possibilities were separately used for the third component concerning the distribution of regulatory costs. One was that all costs were shared equally among individuals regardless of income, a second was that costs were proportional to the income of individuals, and the third was that costs increased with income but at a decreasing rate.

Using the basic model and data from Kitagawa and Hauser and, separately, data from Frerichs et al., the number of fatalities induced by regulations of various costs were estimated. Although these sources used 1960 and 1980 data, respectively, the implications were very similar. In general, when individuals paid equally for regulations, regulatory expenditures between $3 and $3.5 million in 1980 dollars ($4.8 to $5.6 million in 1990 dollars) induced one statistical fatality. When individuals paid for regulations proportional to their incomes, $6 to $7.5 million in 1980 dollars ($9.5 million to $11.9 million in 1990 dollars) induced a statistical fatality. For the situation where individuals with higher incomes paid more for regulations but at a decreasing rate, the cost implications were between the previous estimates.

The model estimates induced fatalities by adding up the very small increases in risks due to reduced income using the relationship in Figure 3 for all the persons, perhaps millions, affected by a regulation. It is important to note that this one relationship is an aggregate one built from statistical information; the relationship between risk and income could be different for each individual and include threshold effects. One might expect that the risk would not increase for most individuals due to a $5 decrease in disposable income. However, the statistics suggest that one person in a million may be at the threshold where the $5 does make a significant difference. In this regard, Figure 3 can be thought of as a relationship that aggregates over all the potential thresholds of many individuals.

The implications of the model were somewhat surprising, suggesting that the cost of regulations may induce many more fatalities than previously recognized. In terms of 1980 dollars, one fatality might be induced for each $3 to $7.5 million cost of the regulation. Because it was clear that such results could be used and misused in evaluating and setting policy, the important caveats of the numbers were clearly mentioned. Still, the results quickly found their way into regulatory agencies, the courts, Congress, and the media.

On July 12, 1991, Judge Williams (1991) of the U.S. Court of Appeals for the District of Columbia ruled against the Occupational Safety and Health Administration (OSHA) concerning a requirement to install locks on energy devices such as circuit breakers. Citing Keeney (1990), he noted in a concurring opinion that a fatality might be induced by each $7.5 million, one reasonable upper bound on the cost of regulation that induces a statistical fatality. If this were the case, the regulation would induce more fatalities than it was expected to prevent.

On March 10, 1992, James B. MacRae, Jr., Acting Administrator and Deputy Administrator of the Office of Information and Regulatory Affairs of the Office of Management and Budget, wrote a letter to the Department of Labor about a proposed OSHA regulation concerning air contaminants in industries. OSHA estimated that the regulation would prevent 8–13 deaths annually, at a cost of $163 million annually. Using the figure of $7.5 million of regulatory expenditures to induce one statistical fatality, MacRae argued that the cost of the regulation could induce 22 additional deaths per year. He suspended review of the draft rule proposed by OSHA and asked OSHA to investigate further the net effect of the proposed regulation on health and safety.

On March 19, 1992, a meeting of the Senate Committee on Governmental Affairs chaired by Senator John Glenn addressed issues raised in MacRae's letter concerning OSHA. MacRae (1992) gave a statement supporting the position of the Office of Management and Budget (OMB). Statements by others tried to refute the concept that richer is safer. In a news release, Senator Glenn (1992) stated, among other things, that OMB has an "Alice-in-Wonderland type claim that health and safety regulations cause harm to workers and the public's health and safety." Glenn further stated that it "seems to stand logic on its head—to say that controlling a dangerous substance in the workplace makes an increased health hazard to the worker." Glenn then asked the General Accounting Office (GAO) to review the proposed OSHA rule and OMB's response. The General Accounting Office (1992) report, which was sent to Glenn on July 2, did not argue with the general model of induced fatalities but said that it was inappropriately applied to the OSHA case. It further stated that the GAO strongly supports "efforts to form more comprehensive and realistic estimates of the effects of government action."

The research on the richer is safer concept and its role in the controversy described above have not escaped the media. Articles that explicitly mention this work are found in *The Wall Street Journal* (Karr and Davis 1992), *Business Week* (Crock 1992), the *Investor's Business Daily* (Merline 1992, Stein 1993), and *Reason* (Marshall 1992). Other articles have appeared in metropolitan newspapers including *The New York Times*, *The Washington Post*, and *The San Francisco Examiner.*

I had four original objectives in pursuing this research. They were to promote broader thinking about the pros and cons of proposed regulation, to change the need for making a value tradeoff between dollars and statistical lives to a need to make a value tradeoff between statistical lives and statistical lives, to provide guidance for the dollars versus statistical lives tradeoff, and to promote communication about the important issue of indirect risks. Due partially to fortuitous circumstances and to the significant research of many others, these objectives have been achieved.

Gathering Information About Potential Nuclear Waste Storage Sites

The Nuclear Waste Policy Act (NWPA) of 1982 outlined the government's plan to dispose of nuclear waste from power plants in a safe and environmentally acceptable manner. The U.S. Department of Energy (DOE) has major responsibility for the plan. This plan involved identification of a site for a nuclear repository, which is a deep underground mine for permanently storing nuclear waste. One step was to identify three sites for simultaneous characterization, which would require numerous activities such as construction of an exploratory shaft, subsurface excavation and tunneling, and testing the host rock in the vicinity of the shaft.

Based on draft environmental assessments produced in December 1984, DOE proposed five sites for nomination as suitable for characterization. These sites (and their host rock types) were in Mississippi (a salt dome), Texas (bedded salt), Utah (bedded salt), Nevada (tuff), and Washington (basalt). With strong prodding from the Board of Radioactive Waste Management of the National Academy of Sciences, DOE chose to conduct a rigorous evaluation of the five sites using multiattribute utility analysis (Merkhofer and Keeney 1987). This analysis, completed in early 1986, indicated the relative desirability of the nominated sites to be Nevada, Mississippi, Texas, Utah, and Washington.

Using insights from this analysis and other information at their disposal, DOE selected the Nevada, Texas, and Washington sites for characterization (DOE 1986). In making this recommendation, DOE necessarily balanced concerns about the diversity of rock types with the evaluations of individual sites. As this balancing was done without the aid of a formal analysis, it had to be based on intuitive professional judgment. The quality of this judgment was called into question by the public and the U.S. Congress. In addition, by 1986, the cost of characterizing a site was estimated to be $1,000 million (i.e., $1 billion) per site. This raised the issue of whether the information expected from characterization of the sites was worth $3 billion. I decided to conduct an analysis to examine these issues: 1) Were the three sites selected the best portfolio of potential repository sites to characterize? 2) Was sequential, rather than simultaneous, characterization of sites a more reasonable way to investigate potential repository sites?

When the NWPA was written, it was recognized that selecting three sites for characterization was a portfolio problem. Because of dependencies between sites (e.g., due to a common rock type), the best three individual sites did not necessarily form the best portfolio. But it was not understood how to appropriately evaluate the diversity of the portfolio. A key observation was that the portfolio chosen and characterized was a means to identify a single final repository site. Characterization was just the gathering of sample information. The second observation was that characterization would provide mainly information about economic costs and little information about the other relevant consequences of site selection, that is, health and safety, environmental, and socioeconomic impacts. Thus, the essential features of the decision could be examined as a sequential decision problem that resolved uncertainties about economic costs only.

Decision trees were used to examine strategies for both simultaneous and sequential characterization. The resulting analyses indicated that selecting the Mississippi site to replace the Texas site in simultaneous characterization would make an improvement equivalent to saving $100 to $400 million. More significant is the fact that sequential characterization strategies were identified that were the equivalent of $1.7 to $2 billion less expensive than DOE's intended simultaneous characterization of the three sites.

The best strategy was to characterize the Nevada site first and select it for the repository, unless its revised equivalent cost estimate was more than $1 billion over the expected cost of the Mississippi site. If this occurred—a 20% chance in the analysis—the Mississippi site should be characterized and the better of the two sites then chosen as the repository.

The analysis (Keeney 1987) was available in late 1986 when Congress began discussions that led to the

Nuclear Waste Policy Amendments Act of 1987. A draft of the amendments stated that, among other things, its purpose was "to provide for the sequential characterization of repository sites [to] result in significant Federal budget savings." The final Amendments Act ended up selecting the Nevada site for the repository.

National Policy to Manage Nuclear Waste

In 1982, when the Nuclear Waste Policy Act was passed, the cost estimates were $35 million for characterization and $1,000 million for site construction and operation. The first repository was to accept spent fuel in 1998. In 1986, when DOE recommended sites in Texas, Washington, and Nevada for characterization, the characterization costs were estimated to be $1,000 million per site, construction and operation had escalated to $8,900 million for the cheapest site, and the startup date was delayed to 2003. The Nuclear Waste Policy Amendments Act of 1987 chose Yucca Mountain, Nevada, as the site for the first repository. By 1989, the estimated cost for that repository was approximately $25,700 million of which $11,500 million were for characterization. The startup date for receiving spent fuel was now 2010.

These cost escalations and delays, coupled with the loss of trust in and credibility of DOE's waste management process, created a virtual impasse for the repository program. Hence, it was appropriate to reconsider the fundamental question: How should the United States manage its nuclear waste generated from power plants? In particular, is it an appropriate use of over $25,000 million to build one repository? Are the benefits of deep geological storage (e.g., providing health and safety protection and safeguarding nuclear materials) worth the loss of retrievability and flexibility in managing the wastes? With von Winterfeldt, we conducted an analysis to help answer these questions by integrating issues of managing nuclear waste into a logical framework.

Strategies to Manage Nuclear Waste. Strategies are the sequences of decisions made from now through 2100 to manage nuclear waste. The consequences of those strategies, and hence their relative desirability, depend both on the decisions taken and on the resolution of uncertainties over time. To clarify the strategies evaluated, it is useful to consider three features: current decisions, resolution of uncertainties over the next 100 years, and decision alternatives in 2100.

One current alternative is storing nuclear waste in an underground repository at Yucca Mountain as soon as possible. The second is an above-ground monitored retrievable storage (MRS) facility to store nuclear waste until 2100. The third is to store nuclear waste above ground at nuclear power plant sites in dry storage casks or concrete bunkers until 2100. All these alternatives would be monitored to ensure that they were safe. Indeed, there is general agreement that the monitored retrievable storage facility and on-site storage can be safely managed for a 100-year period.

As seen in Figure 4, if the initial decision is to pursue the repository at Yucca Mountain, there is uncertainty about whether the facility would be licensed. If Yucca Mountain is not licensed, then another decision must be made about whether to pursue a repository at a new site, a monitored retrievable storage facility, or on-site storage. With either another repository or a monitored retrievable storage

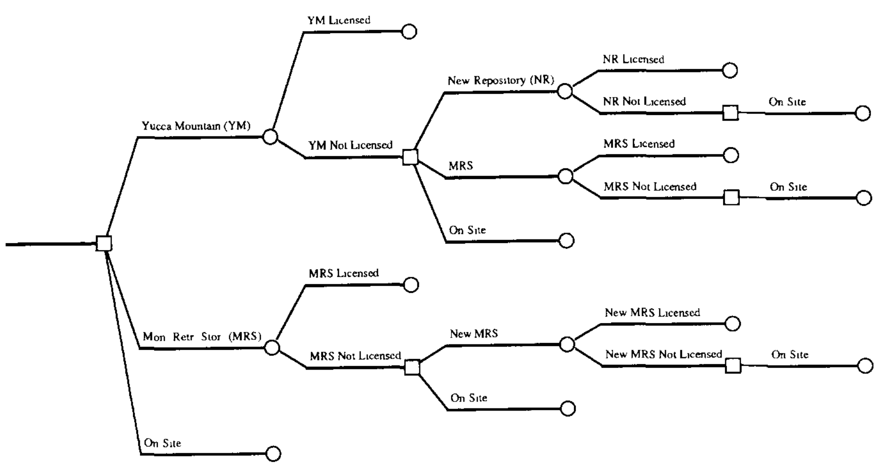

Figure 4. Decision alternatives to manage nuclear waste at the current time (squares denote decisions, circles denote the resolution of uncertainty not under the control of the decisionmakers).

facility, there is again uncertainty about whether it would be licensed.

If the monitored retrievable storage facility is chosen and not licensed, then a decision must be made about whether to pursue such a facility at another site or pursue on-site storage at nuclear power plants until 2100. If a second site for a monitored retrievable storage facility is chosen and not licensed, the nuclear fuel will be stored on site until 2100. In all cases, we assume that on-site storage would be licensed.

Several important uncertainties are resolved after the 11 decision-event sequences in Figure 4. These concern the likelihoods that nuclear material is stolen and/or misused, productive uses are found for some or all of the nuclear waste, there is a cure for cancers induced by radiation, and technological innovations reduce waste management costs.

If a repository exists prior to 2100, the main decision will be whether to dig up the nuclear waste, which depends on its economic value. For situations where monitored retrievable storage or on-site storage have been used through 2100, the desirability of the alternatives depends on uses for the waste, a possible cancer cure, and technological innovation. When all the waste is used, there is no need for a repository. When some of the waste is used, a smaller repository can be built. In addition, when there is a cancer cure, the repository may cost less because of a reduced need to promote absolute safety against potential cancer cases for thousands of years. If technological innovations occur, it will be less expensive to build a repository in 2100. In the case when there is no use for nuclear material, no cancer cure, and no technological innovations, we assume that a repository similar to those that might be built today would be constructed.

Objectives, Consequences, and Value Tradeoffs. The objectives for evaluating nuclear waste management strategies are listed in Table II. The first four objectives (pre- and postclosure health and safety, environment, social impacts) and the seventh (direct cost) have been used in previous studies. In this analysis, four objectives were added: equity of the management strategy in terms of procedure, geographic distribution, and intergenerational burden; fulfillment of the government responsibility to dispose of nuclear waste; indirect economic costs; and consequences of misuse of nuclear waste.

With the problem framed by the objectives, we conducted a decision analysis (Keeney and von Winterfeldt 1994). Attributes were chosen or constructed as indicated in Table II. Information for describing probabilities, such as for licensing, and consequences was gathered from existing data and judgments of persons familiar with the nuclear waste storage problem. We based our selection of the unit value tradeoffs in the table on previous published studies using some of the objectives and on our own judgments augmented by discussions with others. Collectively, all of this provided us with a reasonable set of information for a base case analysis and a very broad sensitivity analysis.

Table II
Objectives, Attributes, and Value Tradeoffs Used for Evaluating Nuclear Waste Management Strategies

Objective	Attribute	Unit Value Tradeoff (Equivalent 1992 Million of Dollars)
1. Preclosure health and safety	X_1: Number of preclosure fatalities due to cancer or accidents	4
2. Postclosure health and safety	X_2: Number of postclosure fatalities due to cancer	1
3. Environmental impacts	X_3: Constructed scale of environmental impacts	80
4. Social impacts	X_4: Constructed scale of social impacts	125
5. Equity	X_5: Component constructed scales	
5.1 Geographic equity	X_{51}: For geographic equity	1,000
5.2 Intergenerational equity	X_{52}: For intergenerational equity	1,000
5.3 Procedural equity	X_{53}: For procedural equity	1,000
6. Government responsibility	X_6: Years until a nuclear waste management system is operational	100
7. Direct economic costs	X_7: Millions of 1992 dollars	1
8. Indirect economic costs	X_8: Component attributes	
8.1 To states and businesses	X_{81}: Millions of 1992 dollars	1
8.2 To electricity ratepayers	X_{82}: Millions of 1992 dollars	1
9. Misuse of nuclear waste	X_9: Constructed scale	10,000

Results of the Analysis. In the base case analysis, the overall equivalent costs were $36,000 million for the repository, $18,000 million for the monitored retrievable storage facility, and $8,000 million for on-site storage. Hence, the best strategy for managing nuclear waste is to pursue on-site storage until 2100. At that time, no decision is necessary if all of the waste is useful. If all of the waste is not useful, a repository should be built at a cost consistent with the conditions of that time. Numerous sensitivity analyses supported the base case conclusions.

Not placing nuclear waste in a repository at this time affords us flexibility for future decisions. That flexibility would enable us to take advantage of any future use for nuclear waste, potential cancer cure, and technological innovations for managing the waste. The value of flexibility, calculated as the expected economic savings due to these factors, is about $11,000 million. Even with no waste use or cancer cure, the value of flexibility is still $6,600 million.

One major reason to build a repository at Yucca Mountain is to isolate the radioactivity from future generations living thousands of years from now. However, the postclosure health effects from any possible releases of radiation material are identical for all three strategies, because all three end up either with no waste because it is all used by 2100, or with a repository. Thus, potential postclosure fatalities (often stated as the reason why a repository is needed soon) cannot be used to differentiate the desirability of the strategies.

An alternative that could be implemented now is to create a fund for society in 2100 to determine how it wishes to dispose permanently of nuclear waste from our current power plants. For this analysis we assumed that spent nuclear fuel will be stored on site until 2100. We based the size of a potential fund on the historical long-term real interest rate, which depends on the inflation rate and actual return rate. Using data from 1926–1991, real interest growth rate on long-term government bonds, a very conservative investment, was 1.35% annually. Compounded at this rate for 108 years until 2100, one dollar grows to 4.26 real dollars. Thus, a $6,000 million fund established today would grow to $25,700 million by 2100, which is just enough to finance a repository at today's prices.

Using base case assumptions, the on-site storage alternative with a $6,000 fund is preferable to the on-site storage alternative without a fund by $31,000 million, to the repository without a fund by $59,000 million, and to monitored retrievable storage without a fund by $40,000 million. Since the entire current cost of a repository built in 2100 is covered by a $6,000 million fund, the expected cost savings of $13,200 million due to possible technology innovation, waste use, and/or a cancer cure is due to the fund alternative. Most of the remaining savings are due to changes in the indirect costs eventually borne by electricity ratepayers. However, even when ignoring this attribute, the savings are still $19,000 million compared to on-site storage without a fund, higher when compared to the repository and the MRS without a fund.

The insights from this analysis suggest three major recommendations: stop pursuing Yucca Mountain as a repository to permanently store nuclear waste, modify the Nuclear Waste Policy Act to allow the consideration of a broader range of strategies to manage nuclear waste, and conduct an independent detailed analysis along the lines of the analysis outlined here to help design a better national policy.

The simple fact is that our information today is much different from that available at the time the Nuclear Waste Policy Act was passed that specified as our national policy for storing nuclear waste a permanent underground repository built as soon as possible. Partially as a result, the current attempt to store nuclear waste in a repository at Yucca Mountain appears to be a major misallocation of national resources.

Common Aspects of These Problems

Legislators, regulators, stakeholders, and the public care about each of these four problems because of their values. The alternatives chosen for these problems could have large consequences for the achievement of these values. The relevant values were explicitly included in the analyses for review, sensitivity analyses, and approval. They were also included in a way that can be communicated to and understood by interested individuals who have no analytical training. I believe that including of the values of interest in this way is the main reason why the analyses seemed to affect subsequent thinking and policy.

What was it that the four problems addressed had in common? They were important in that their potential consequences were significant, they were messy, implying that it was not obvious what needed to be done, and they were understudied by individuals with the knowledge of operations researchers and management scientists. All these problems were examined from many perspectives by lawyers, economists, political scientists, sociologists, politicians, and lobbyists. But their approaches and resulting insights differ from ours. We in operations research have some unique contributions to make in understanding such problems. We think differently. We have the training to address complexities inherent in such problems—complexities like multiple objectives, uncertainties, the lack of data and/or conflicting data, and the need for value judgments.

ADDRESSING MEGA-RISKS TO OUR COUNTRY

Let me offer a challenge to you. Consider your professional objectives carefully. Then see whether a useful way to pursue them is to focus some of your talent on what I consider to be the mega-risks facing our country.

The United States is a bit over 200 years old. What is the likelihood that she will still exist, with roughly the same or a better quality of life and the same economic, political, and geographic position in the world, on her 300th birthday? Over the past few years, I have asked many people this question. Nobody has answered 100%; nobody has answered 95%. The range of responses is from 25 to 90%. So perhaps we should examine the major risks that threaten our country. Perhaps we should study and estimate these risks. Perhaps we should develop and evaluate policy alternatives that would reduce them.

To examine risks to America—what I think of as mega-risks—we would first need to identify them. Some important contenders are the following: lawyers and their role in our society, shortcomings of our educational system, illegal immigration, use of drugs, breakdown of the family, unwanted and uncared-for children, abortion, violence, gun control, competitiveness of American industry, white collar crime, government regulation, and government gridlock. Many of these sources of risk are intertwined.

The systematic evaluation of risks to America requires that the values that the country represents be clearly articulated. In simple terms, these values may be to provide for life, liberty, and the pursuit of happiness. Then the risks threatening the country are those that impinge upon the achievement of these values. To investigate this, we must articulate in more detail the meaning of life, liberty, and the pursuit of happiness.

With the values clearly stated, we can begin to estimate the risks. This task can be facilitated using models based on logical reasoning and obtainable information. In essence, we need to relate current policies of our country to the resulting impacts on the achievement of our country's values. Although the development of such a model can be an extreme challenge, I believe the potential insights and eventual effect on society are worth the effort. It may be helpful to describe a couple of examples. Let me be specific, but simplistic for illustration.

Consider illegal immigration. Illegal immigration and legal immigration are not the same. Both have many consequences to America, some good, some bad. Operations researchers can do analyses that clarify the distinctions by estimating those consequences. If an analysis indicated that the net implication to our country is positive from legal immigration and negative from illegal immigration, that insight may be useful. If an analysis suggests that both legal and illegal immigration had net positive effects, it might be reasonable to examine a potential policy that cut illegal immigration and increased legal immigration. This may have greater positive effects.

Analyses addressing illegal immigration should consider the full range of consequences. Illegal immigration results in the influx of many more persons, which affects the quality of life in many ways, from traffic congestion to air pollution to crime. It taxes medical and educational systems, reducing medical and educational resources for Americans and legal immigrants. Information on illegal immigrants today suggests that they are not the same as legal immigrants in earlier periods; they are on average prepared to work only in the lowest-level jobs. This eliminates job opportunities for Americans with little education, while it also keeps prices of certain products, such as farm products, lower. In California alone, the annual cost of illegal alien felons in prisons is approximately $500 million. Welfare and criminal justice system costs are also higher. The large number of illegal aliens makes it harder to have a single language in our country, a potential bond for holding together people from different countries, races, and religions. The multiple languages in many countries (such as Canada, Belgium, and the former Soviet Union) cause or exacerbate many serious problems. How do these problems affect the risk to our country's existence? Not an easy question to answer, but certainly one worthy of much deep and structured thought. Examining policy options that might reduce these risks is also worthy of serious effort.

Consider a different type of risk, lawyers and their role in society. Lawyers make the laws: They are the majority of legislators. And lawyers interpret the laws: They are the judges. They are the only group in America not subject to checks or balances by some different group. Estimates of lawyers' annual fees have been in the range of $300 billion. Each year there is a transfer of substantial amounts of money from nonlawyers to lawyers to pay for this. Nonlawyers receive benefits from lawyers' efforts, but they also become economically poorer. Being poorer results in a lower quality of life and in the additional loss of life, as poor people have a higher death rate than those with more disposable income. Most lawyers are bright, so there is a brain drain into the legal field. Our country loses relative to the circumstance where these bright people go into fields that produce goods or services of greater value. In general, business people and engineers increase the pie, lawyers reduce and divide it. Laws and the legal process hold up production and innovation, and in some cases stop production and services because of suits or threatened suits. Jobs are moved to other countries to avoid the potential litigation that is so much a part of our society. This certainly puts a dent in the pursuit of happiness for many Americans.

It would be very useful to develop a model to estimate the various pros and cons of the way in which lawyers now operate in America. The analysis would need to distinguish between different types of lawyers. It would need to examine indirect as well as direct ways in which the quality of life in America and the future existence of America are influenced. The analysis should develop and appraise alternative policies that would alter the role of lawyers in our society. The result should be to focus thinking on several critical issues, raise the quality of communication about these issues, and influence, perhaps, the course of affairs.

In summary, the challenge is to frame some of the critical problems facing our nation by using our unique knowledge as operations researchers. This framing will provide important insights for politicians, the media, and citizens. It should provide a framework to enhance productive discussion of these issues. More important, it will guide thought and action on absolutely crucial problems facing us. If done well, the analyses and their insights should justify the necessary time and effort.

ACKNOWLEDGMENT

It is a tremendous honor to receive the Philip McCord Morse Lectureship Award. I have many people to thank for this award. The roles that Professors Morse, Drake, and Raiffa played in guiding my career were critical. I have been extremely fortunate to work with each of them.

My colleagues have enhanced my career and my life. They are the source of ideas, enthusiasm, and enjoyment. I would like to recognize the contributions in this regard of my colleagues as graduate students at MIT, as faculty at MIT, and the University of Southern California, as research scholars at the International Institute for Applied Systems Analyses in Austria, and as scientists and engineers at Woodward-Clyde Consultants. For the past twenty years, I have had the pleasure of working with Professor Detlof von Winterfeldt of USC on numerous research and consulting projects. With him, these projects are always both interesting and fun.

Finally, this paper was written with partial support from grant SBR-9308660 from the National Science Foundation.

REFERENCES

BUREAU OF THE CENSUS. 1981. Statistical Abstract of the United States. Department of Commerce, Washington, D.C.

CROCK, S. 1992. What's Deadlier, A Toxic Workplace or Low Pay? *Bus. Week*, April 13, 100–101.

DEPARTMENT OF ENERGY. 1986. Recommendation by the Secretary of Energy of Candidate Sites for Site Characterization for the First Radioactive-Waste Repository. DOE/S-0048, Office of Civilian Radioactive Waste Management, Washington, D.C.

DRAKE, A. W., R. L. KEENEY AND P. M. MORSE. (EDS.). 1972. *Analysis of Public Systems*. MIT Press, Cambridge, Mass.

DYER, J. S., AND R. K. SARIN. 1979. Measurable Multiattribute Value Functions. *Opns. Res.* **27,** 810–822.

ENVIRONMENTAL PROTECTION AGENCY. 1983. Background Information Documents: Proposed Standards for Radionuclides. EPA 520/1-83-001, Office of Radiation Programs, Washington, D.C.

FISHBURN, P. C. 1965. Independence in Utility Theory With Whole Product Sets. *Opns. Res.* **13,** 28–43.

FRANKLIN, B. 1772. Letter to Joseph Priestly. Reprinted in *The Benjamin Franklin Sampler*, 1956. Fawcett, New York.

FRERICHS, R. R., J. M. CHAPMAN, P. NOURJAH AND E. F. MAES. 1984. Cardiovascular Diseases in Los Angeles, 1979–1981. American Heart Association, Greater Los Angeles Affiliate, Inc., Los Angeles.

GENERAL ACCOUNTING OFFICE. 1992. Risk-Risk Analysis, OMB's Review of a Proposed OSHA Rule. Report of the Chairman, Committee on Governmental Affairs, U.S. Senate, Washington, D.C. (July).

GLENN, J. Senator 1992. News Release. March 19.

GREGORY, R., AND R. L. KEENEY. 1994. Creating Policy Alternatives Using Stakeholder Values. *Mgmt. Sci.* (to appear).

KARR, A. R., AND B. DAVIS. 1992. Bush to Require Regulators to Weigh Costs and Impact on Health, Mortality. *The Wall Street Journal*, March 20, p. A3.

KEENEY, R. L. 1968. Quasi-Separable Utility Functions. *Naval Res. Logist. Quart.*, December, 551–565.

KEENEY, R. L. 1974. Multiplicative Utility Functions. *Opns. Res.* **22,** 22–34.

KEENEY, R. L. 1987. An Analysis of the Portfolio of Sites to Characterize for Selecting a Nuclear Repository. *Risk Anal.* **7,** 195–218.

KEENEY, R. L. 1990. Mortality Risks Induced by Economic Expenditures. *Risk Anal.* **10,** 147–159.

KEENEY, R. L. 1992. *Value-Focused Thinking: A Path to Creative Decisionmaking.* Harvard University Press, Cambridge, Mass.

KEENEY, R. L. 1994. Creativity in Decision Making With Value-Focused Thinking. *Sloan Mgmt. Rev.* Summer, 33–41.

KEENEY, R. L., AND T. L. MCDANIELS. 1992. Value-Focused Thinking About Strategic Decisions at BC Hydro. *Interfaces* **22,** 94–109.

KEENEY, R. L., AND H. RAIFFA. 1993. *Decisions With Multiple Objectives.* Cambridge University Press, New York. (First published in 1976 by Wiley, New York).

KEENEY, R. L., AND D. VON WINTERFELDT. 1986. Why Indirect Health Risks of Regulations Should Be Examined. *Interfaces* **16,** 13–27.

KEENEY, R. L., AND D. VON WINTERFELDT. 1994. Managing Nuclear Waste From Power Plants. *Risk Anal.* **14,** 107–130.

KITAGAWA, E. M., AND P. M. HAUSER. 1973. *Differential Mortality in the United States of America: A Study of Socioeconomic Epidemiology.* Harvard University Press, Cambridge, Mass.

LUTTER, R., AND J. MORRALL. 1994. Health-Health Analysis: A New Way to Evaluate Health and Safety Regulation. *J. Risk and Uncert.* **8,** 43–66.

MACRAE, J. B., JR. 1992. Statement Before the Senate Committee on Governmental Affairs, March 19, Washington, D.C.

MARSHALL, J. 1992. Risky Business. *Reason*, August/September, 52–53.

MERKHOFER, M. L., AND R. L. KEENEY. 1987. A Multiattribute Utility Analysis of Alternative Sites for the Disposal of Nuclear Waste. *Risk Anal.* **7,** 173–194.

MERLINE, J. 1992. Is OSHA Really Protecting Us? *Invest. Bus. Daily*, August 6, 1–2.

MORSE, P. M., AND G. E. KIMBALL. 1951. *Methods of Operations Research.* MIT Press, Cambridge, Mass. and John Wiley, New York.

NADLER, G., AND S. HIBINO. 1990. *Breakthrough Thinking.* Prima Publishing and Communications, Rocklin, California.

PETERS, T. J., AND R. H. WATERMAN JR. 1982. *In Search of Excellence.* Harper & Row, New York.

PRATT, J. W., H. RAIFFA, AND R. O. SCHLAIFER. 1964. The Foundations of Decision Under Uncertainty: An Elementary Exposition. *J. Am. Statist. Assoc.* **59,** 353–375.

RUSSO, J. E., AND P. J. H. SCHOEMAKER. 1989. *Decision Traps.* Doubleday, New York.

STEIN, R. S. 1993. What Cost Environmentalism? *Invest. Bus. Daily* July 13, 1–2.

VON WINTERFELDT, D., AND W. EDWARDS. 1986. *Decision Analysis and Behavioral Research.* Cambridge University Press, New York.

WILDAVSKY, A. 1979. No Risk is the Highest Risk of All. *Am. Scient.* **67,** 32–37.

WILDAVSKY, A. 1980. Richer is Safer. *The Public Interest.* **60,** 23–39.

WILDAVSKY, A. 1988. *Searching for Safety.* Transaction Books, New Brunswick, New Jersey.

WILLIAMS, JUDGE. 1991. *UAW vs. OSHA.* United States Court of Appeals, District of Columbia, July 12.