

13 Practical Value Models

Ralph L. Keeney and Detlof von Winterfeldt

ABSTRACT. Many complex decision problems require value judgments, and it is often useful to build formal models of these value judgments. Several models and assessment procedures exist for this task ranging from simple rating and weighting techniques to sophisticated multiattribute utility models that incorporate tradeoffs and risk attitudes. In this chapter, we argue that the choice of a value model and assessment procedure depends on the decision context, the purpose of the model, the quality of available value judgments, time, and other factors. We propose the notion of a *practical value model* that is useful for a given decision context and purpose, using simplifications and approximations, when appropriate. We then discuss practical considerations when choosing a value model, for example, when deciding whether a value model is additive over attributes or not. We also consider practical issues when choosing an assessment procedure, for example, when selecting equivalent-cost procedures versus rating and weighting procedures. We discuss several topics related to interacting with decision makers when constructing value models, and we conclude with the description of two generic practical value models that we have found useful in many past applications.

Two types of models are commonly used in decision analysis: *consequence models* incorporate the facts, judgments, and uncertainties inherent in decision problems to describe or predict possible consequences of alternatives and *value models* incorporate the values or value tradeoffs and risk tolerances to evaluate consequences. This chapter addresses the value modeling aspect of decision analysis. It focuses on the choices that a decision analyst has to make when selecting an appropriate value model and assessment procedure.

Most people are familiar with value models, but they may not realize it, as the term is somewhat new and not in common use (Keeney 1988). In operations research and management science, one term for a value model is an *objective function*, often referring to a single objective, such as maximizing profit or minimizing delays. In decision analysis, value models are sometimes called utility functions, value functions, or preference functions. We prefer to use the general term *value model*, as the process of constructing such a model is the same as the process used to construct any other type of model. The analyst decides on the variables to use in constructing the value model, selects and/or verifies the relationships among those variables to get a mathematical representation of interest (preferences in this case), and then quantifies the parameters for the model using information that is available and/or can be gathered.

A value model is built such that a number, referred to as a value, can be calculated for each of the alternatives considered in a decision. These numbers provide a ranking of the desirability of each alternative, thereby providing the basis for making a choice. In many cases we know of, the value models involve multiple conflicting objectives. In some cases the value model provides a ranking of desirability and in some cases of undesirability, suggesting either to maximize or to minimize the value.

This chapter does not present the mathematical theories for constructing value models. There are many references for these theories (e.g., Krantz et al. 1971; Keeney and Raiffa 1976; and von Winterfeldt and Edwards 1986). This chapter is about the appropriate use of these theories for building a value model. Specifically, we focus on how one should build a value model that is appropriate and practical for a specific decision problem. We call such a model a “practical value model.” We also are concerned with the practical assessment procedures that go along with a choice of a practical value model.

We add the word “practical” to suggest that it is not always necessary or useful to construct a state-of-the-art value model that is completely justified on theoretical grounds or to use theoretically justifiable assessment procedures. Approximations may do well both in modeling and assessment. The choice of value models and assessment procedures is a function of the characteristics of the decision being faced, the characteristics of the decision maker or makers, the time available for the process, and the skills of the decision analyst who is facilitating that process. A practical value model needs to balance the theoretical soundness of the model to represent the details of the decision makers’ preference, the time and effort needed to construct the model, the quality of insight that can be gained from using the model, and the ability of decision makers to understand the insights from the model. An implication is that a practical value model should be “good enough” so that it is not the weak link in the overall analysis of the decision. If the refinement of the alternatives or the quality of the information about how well the alternatives measure up in terms of the objectives is such that additional legitimate insights could not be gained by an analysis with a more sophisticated or “better” value model, then the value model is good enough in terms of the overall analysis. If the analysis is of a quality such that the main limits to its use are the willingness of the decision makers to use it, or their ability to understand the insights, then the value model is good enough given the process.

Building a Value Model

There are two types of value models: *Utility functions* and *value functions*. The theoretical foundations of utility functions are found in von Neumann and Morgenstern (1947) and Keeney and Raiffa (1976). The theoretical foundation for value functions are found in Krantz et al. (1971) and Dyer and Sarin (1979). Krantz et al. develop value functions based on the notion of “strength of preferences.” Dyer and Sarin refer to value functions as “measurable value functions.” Constructs used for the existence of utility and value functions in these references

are our basis for judging the appropriateness or theoretical correctness of value models. We can simplify these models using assumptions for a specific decision problem that are practical as described above.

In a given decision problem, analysts have to choose among a variety of possible value models. In essence, the question is what value model should be built for this decision? Some notation will be useful. We will assume that a set of alternatives $A_j, j=1, \dots, J$ has been identified for the decision problem. Also, suppose we have identified a set of fundamental objectives $O_i, i=1, \dots, n$ using concepts such as those described in Chapter 7 in this volume. Also, assume that an appropriate attribute $X_i, i=1, \dots, n$ has been defined respectively to measure each of those objectives. For instance, in a public policy decision about highway safety, one fundamental objective would be to “minimize the loss of life in traffic accidents,” and a corresponding attribute might be the “number of individuals dying in traffic annually.” With this notation, a possible consequence of any chosen alternative can be written as (x_1, \dots, x_n) or more compactly as \mathbf{x} , where x_i is a specific level of X_i .

Once a set of fundamental objectives and associated attributes are defined, the next step in constructing a value model is to obtain a general structure for the preference relationship among the various \mathbf{x} s. Building a value model is analogous to building a factual model of, for instance, automobile accidents that might relate driver behavior, vehicle features, weather conditions, and road quality to subsequent accidents and highway fatalities. To combine the attributes in a value model, one uses independence concepts analogous to probabilistic independence or conditional probabilistic independence used in factual models. The main types of independence concepts are *utility independence*, *preferential independence*, and *additive independence*. Descriptions of these independence concepts and the implications for the functional forms of the resulting value models are described in detail in Keeney and Raiffa (1976) and von Winterfeldt and Edwards (1986).

It is also important in a value model to describe the preferences for consequences that differ in terms of only a single attribute. The common concepts for single-attribute utility functions concerned with uncertainties of consequences are risk aversion, risk neutrality, and risk proneness. For the measurable value functions analogous conditions are marginally decreasing, constant, or increasing value.

Certain sets of assumptions about preferences among consequences imply a specific functional form of the value model. An example is the additive value function

$$v(x_1, \dots, x_n) = \sum w_i v_i(x_i), \quad (13.1)$$

where v is the overall value model, the v_i are single attribute value functions, and the w_i are scaling constants. The original work for this type of value model was presented in Luce and Tukey (1964) and Krantz (1964), who considered only ordinal judgments when constructing this value function. Fishburn (1970) and Krantz et al. (1971) considered value functions based on strength of preference judgments.

When alternatives involve uncertainty, it is common to begin with a utility function and, depending on appropriate assumptions, develop decompositions over attributes. The most common decompositions are the additive and multiplicative ones (see Keeney and Raiffa 1976):

$$u(x_1, \dots, x_i, \dots, x_n) = \sum k_i u_i(x_i), \quad (13.2)$$

or

$$1 + ku(x_1, \dots, x_n) = \pi \{1 + k k_i u_i(x_i)\}, \quad (13.3)$$

where u is the overall value model, the u_i are single-attribute utility functions, and the k_i are scaling constants. The seminal work on assumptions for an additive utility function is Fishburn (1965).

The v_i , w_i , u_i , k_i , and k terms in Eqs. 13.1 – 13.3 are the necessary ingredients to construct the value model. These must be determined using value judgments from the decision maker or assigned to represent the value judgments of the decision maker. These value judgments are the “data points” that one needs to construct the value model. For such value data, the basic information exists in the minds of the decision makers or individuals knowledgeable about that decision maker’s preferences and the given problem. This is different from collecting data to assign parameters in a consequence model of automobile accidents, for instance. For consequence models the data is “out there” in the world and can be observed or inferred by analyzing drivers, vehicles, road surfaces, weather conditions, and accident information.

The value data to assign parameters is typically gathered by finding different consequences that are of equal value to the decision maker. Then, using a value model such as Eq. 13.1, the values of the two different consequences, which we will call y and z , can be equated using Eq. 13.1 to yield:

$$v(y) = \sum w_i v_i(y_i) = \sum w_i v_i(z_i) = v(z), \quad (13.4)$$

which provides one equation with the w_i s as the n unknown parameters to be specified. By finding n pairs of consequences of equal value, we generate n equations with collectively n unknowns and then solve for those unknowns, which gives us numerical values of the w_i parameters.

Practical Considerations for Choosing a Value Model

A value model must explicitly address the complex and important value judgments that can influence a decision. If fun and excitement are important objectives of a major vacation, then those objectives should be explicitly included in a value model to evaluate alternative vacation destinations. If one of the main impacts of a decision is the potential loss of life, such as decisions involving the setting of national ambient air quality standards, then an objective should explicitly address the potential loss of life. The objective should be “minimize loss of life” with a corresponding attribute such as “number of lives lost annually because

of the pollutant.” Analyses that use an attribute such as parts per million of the pollutant in the air as a proxy for loss of life are inadequate and miss the substance of the decision. Anybody using such an analysis would have to informally relate the parts per million of the pollutant to the loss of life to make a reasoned decision.

For decisions involving the potential loss of life, value tradeoffs between cost and potential loss of life are typically critical to the choice of an alternative. In those cases, that value tradeoff should be included in the value model. In an investment decision, the risk attitude concerning financial outcomes may be very important and, if that is the case, should be carefully modeled.

In general, whenever a value judgment is important to a decision, it should be explicitly included in the value model. In a personal decision about whether to have amniocentesis or chorionic villus sampling during pregnancy to assess any potential genetic damage of the fetus, a reasoned decision partly depends on the relative desirability to the parents of having a perfectly healthy child, no child, or a child with the various congenital problems caused by gene defects. For this decision, the relative desirabilities of these different situations are sometimes more crucial to making an appropriate choice for a family than are the exact probabilities of the different outcomes.

Utility versus Measurable Value

An analyst can choose whether or not to develop a utility function or a measurable value function. Which of these two functions is appropriate depends on many considerations. Theoretically, the choice mainly depends on the problem characteristics. If uncertainties are large and explicitly incorporated in the consequence model for the decision, then a utility function is the theoretically appropriate value model. If there are no uncertainties or only minor uncertainties about the consequences, then a measurable value function is more appropriate.

Beyond the theory, there are practical considerations that are relevant. If the previously mentioned risk aversion of one of the objectives was important, then it may matter and a utility function would be appropriate. However, if there are only minor uncertainties and if the degree of risk aversion is small, then it probably does not matter whether one chooses a value or a utility function. In such cases, a more important consideration is the relative ability of the decision maker to provide information about values during the assessment of value or utility functions. Generally speaking, the assessment of value functions is easier, involving primarily ordinal or strength-of-preference judgments between pairs of consequences. The theoretically correct assessment of utility functions involves preferences among lotteries or uncertain prospects and some decision makers have difficulty expressing preferences among these often complex options.

One very important result of the research on value and utility functions is that they must be closely related. In all cases, value and utility functions must be related by a strictly monotone increasing function h :

$$u(\mathbf{x}) = h\{v(\mathbf{x})\}. \quad (13.5)$$

Further restrictions on the function h obtain when both the value function and the utility functions are additive. In that case the function h must be a linear positive transformation (see Krantz et al. 1971), so

$$u(\mathbf{x}) = av(\mathbf{x}) + b, a > 0, \quad (13.6)$$

where a and b are constants.

Another example is the case in which the value function is additive, but the utility function is multiplicative. In this case, the value and utility functions must be related by an exponential function (see Dyer and Sarin 1979; Keeney and Raiffa 1976; Barron, von Winterfeldt and Fischer 1984), so

$$u(\mathbf{x}) = a \exp^{cv(\mathbf{x})} + b, a > 0, \quad (13.7)$$

where a , b , and c are constant. As a byproduct of these functional relationships between u and v , the relationships between the single-attribute functions v_i and u_i are also highly constrained to be either linear or exponential. Proofs of these relationships can be found in Keeney and Raiffa (1976), Dyer and Sarin (1979), and von Winterfeldt (1979).

If it can be established that a value function is additive and that the utility function is either additive or multiplicative, a sound way to proceed is to construct an additive value function, followed by an assessment of the parameters of the function h in Eq. 13.5. This function is sometimes called the “relative risk aversion function” because it is the risk aversion expressed relative to a value function.

Some research suggests that utility functions and measurable value functions should be the same (Bell and Raiffa 1988). Other empirical research indicates that people distinguish between these concepts in some situations. For example, Barron, von Winterfeldt, and Fischer (1984) showed that empirically assessed utility functions differ from value functions in the sense of showing relative risk aversion. An exponential relationship between value and utility provided a good fit to their data.

For many decisions, it is reasonable to assume that the measurable value function and the utility function are identical. It is often the case that the quality of the information describing the consequences of the alternatives is such that this assumption about a common function for utility and measurable value is not the weak link in any analysis. In a book written for individuals to help themselves analyze decisions (Hammond et al. 1999), the assumption was made that a utility function and a measurable value function are identical.

Additive versus Nonadditive Value Models

An additive value model is simpler to construct and to use than a nonadditive model. The key to determine if an additive value model is appropriate is the set of objectives used for the value model. If these objectives are fundamental objectives for the decision and meet the criteria specified for a good set of fundamental objectives (see Chapter 7 in this volume), then a strong case can be made that an additive value model is appropriate. In most of our own recent applications of

decision analysis, we have created a set of fundamental objectives that satisfied these criteria so the resulting value model was additive.

The main cause of nonadditivity is the use of means objectives in the value model (see also Chapter 7 in this volume for the distinction between *means objectives* and *ends objectives*). Suppose one is concerned about the loss of life from automobile accidents. One means objective is to minimize speeding of drivers, and a second means objective is to minimize driving under the influence of alcohol. Any value model that simply added up the two effects related to each of these objectives would not incorporate the natural dependency that speeding while under the influence of alcohol is likely worse than the sum of the consequences of each separately. Another case of nonadditivity occurs, when one objective works as a multiplier on a second objective. For example, when evaluating a new piece of equipment, its reliability and its performance may be objectives. In this case reliability acts like a multiplier, because with zero reliability, one should not care about the degree of performance. If one encounters a case such as this, it is often better to develop another objective such as “expected performance” and to develop a model that combines reliability and performance to measure this new objective.

Formally, Fishburn’s (1965) necessary and sufficient condition for the additive utility function (13.2) was that preferences only depend on the marginal probability distributions over consequences. Dyer and Sarin (1979) showed that a measurable utility function is additive, if the strength of preference between two alternatives that vary only on a subset of attributes remains the same, if the consequences in the other attributes that are identical across alternatives are changed.

It is worth mentioning that many value models may have several additive terms and perhaps one or two nonadditive terms. However, even in these situations one might interpret the value model as being additive. To illustrate this, consider a simple decision problem where a decision maker wishes to please two separate individuals A and B. the decision maker’s objectives are explicitly stated as “maximize the desirability of the consequences to individual A” and “maximize the desirability of the consequences to individual B.” If there are major uncertainties, the value model should be a utility function. In these situations, it is reasonable that utility independence assumptions hold so that the decision maker’s utility function u can be written as

$$u(u_A, u_B) = k_A u_A + k_B u_B + k u_{AB}, \quad (13.8)$$

where u_A and u_B are utility functions representing preferences of individuals A and B and the k s are constants. If k is not 0, then this function is commonly thought of as a nonadditive value model. This thinking is based on a value model with two objectives and three terms, one of which is a product term. On the other hand, Eq. 13.8 is clearly additive with three terms. The first two terms directly measure the degree to which the two stated fundamental objectives are met. But what does the third term measure? As discussed in Keeney (1981), the third term can be considered a measure for the objective of the decision maker “to have an equitable consequence for individuals A and B.” If this third objective is added to

the other two fundamental objectives, then Eq. 13.8 is clearly an additive utility function for three fundamental objectives.

There are two important points of this simple example. First, it emphasizes that with an appropriate set of fundamental objectives, an additive value model is practical and logically sound. Second, it indicates that one can get most of the benefits of an additive value model, which is simpler than nonadditive models, and, yet, incorporate a term that may not be additive to capture important consequences in the decision. Even if one does not identify an appropriate label to characterize the nonadditive term with a new fundamental objective, the decision maker should understand why it is included in the value model and how it can easily be handled in practice.

Linearity

Three general situations lead to linearity of single-attribute value or utility functions. One situation is when the attribute for a fundamental objective measures something that is of value in itself, as opposed to value for its uses. Suppose one is evaluating several medical programs and one of the objectives is to save as many lives as possible. The attribute for this objective might be the number of lives saved. Clearly each life saved is valuable in and of itself. It is of course valued by the individual whose life is at risk and by their family and friends. It seems reasonable to count each life saved as equal to any other, especially when the medical intervention is typically used for individuals of roughly the same age and the same health state.

How could one justify a value function, which valued the first life saved much more than the hundredth life saved? One reason might be that the first life saved proved the feasibility of the intervention being used. Subsequently, this may lead to more lives being saved in the future. We certainly recognize that proving feasibility of an intervention, and saving more lives in the future are important, but we feel these should be identified as separate objectives for the decision problem. That being the case, these reasons could not be any justification for valuing the first life save more than later lives saved. The point is that if the value of saving lives is basically for those whose lives are saved and those who know them, counting each life saved equally is reasonable and the implication is that a linear value model for that single attribute is reasonable.

A second situation that leads to linearity is the following. Sometimes the consequence estimates are simply estimates thought to be somewhere around the expected consequence. If one could quantify a reasonable standard deviation in this case, it may be 20 percent or more of that expected consequence. By using the expected value, linearity is naturally implicitly assumed. Therefore, it does not make any sense to use a value model other than a linear model in such situations.

The third situation where linearity is a reasonable assumption concerns single decisions when many other decisions also contribute to the same objective. Therefore, the range over which the consequences may fall in that single decision is small relative to the range for all the decisions collectively. A specific case is

when a major company that has profits in the billions of dollars per year is making a decision that may at most affect annual profits by plus or minus \$20 million. Any reasonable utility function over profits for the company, whether it was risk averse or risk prone, would effectively be equivalent to a linear function over a \$40 million range. Any additional benefit of using a nonlinear utility function would likely not be worth the shortcomings of additional effort for such a decision.

In some circumstances, the linearity assumption may be reasonable for 99 percent of the consequences evaluated in a decision problem, but there may be one or two consequences where there should be an adjustment because linearity may not be appropriate. For instance, a recent study involved a major project in the area of a large dam (Keeney and von Winterfeldt, 2005). The project cost was in the neighborhood of \$600 million and, as a result of the project, there will be a very small chance that there could be a dam failure induced. Alternatives were available that would reduce that small chance and cost up to \$26 million. There were other relatively minor problems that could be caused in the area of the dam such as landslides in local neighborhoods. Considering all the consequences in this decision problem and the cost of the alternatives considered, the range of consequences was about \$30 million, except for the dam failure, which was valued at \$1.2 billion. In this situation for reasons discussed in the previous paragraph, a linear model was appropriate for all possible consequences except those involving a dam failure. In this case, we tripled the consequences of a dam failure to \$3.6 billion, which effectively is the same as utilizing risk aversion for evaluating the consequence of dam failure.

Figure 13.1 illustrates this situation. All of the consequences except the dam failure are in region A and the risk averse utility function over costs shown in the figure is very close to linear in that region. A \$1.8 billion consequence, which is \$0.6 billion for construction and \$1.2 billion due to a dam failure, evaluated at point B with a risk averse utility function, has the same utility as the \$4.2 billion consequence, which is \$0.6 billion plus a dam failure cost tripled to \$3.6 billion, evaluated at point C, using a linear utility function.

Risk Aversion

In some decisions, it is reasonable to assume that the utility function or the measurable value function for single attributes is not linear. If these are utility functions, the main issue is whether they should exhibit risk aversion or risk proneness. It is usually practical in these situations to select utility functions that are appropriately risk prone or risk averse that have one parameter. Specifically, based on the work of Pratt (1964), it is often appropriate to select the constantly risk-averse exponential utility function:

$$u_i(x_i) = a_i + b_i(-e^{-c_i x_i}), b_i, c_i > 0, \quad (13.9)$$

where a_i , b_i , and c_i are constants. With an exponential utility function, it is easy to conduct sensitivity analyses to determine if either the decision or the insight from




Figure 13.1. Adjustment of consequences to allow use of linearity assumption for preferences.

the analysis changes as a result of varying the risk coefficient over a reasonable range.

As mentioned in the discussion about the relationships between utility and measurable value, one can develop a measurable value function for use in a decision problem and use this as a single attribute in a utility function if uncertainties are significant. One can also use the exponential utility function in this case and conduct sensitivity analyses to examine the relevance of risk aversion to the decision.

There are some situations where it is appropriate to use a risk-averse utility function that may be more sophisticated than assuming constant risk aversion. The main situations where this occurs are in investment decisions where the range of financial consequences extends over significant amounts of money and/or long time periods. For individuals doing financial planning for their lifetime, it may be reasonable to use a decreasing risk-averse utility function that automatically accounts for the wealth level of the individual. Meyer (1970) used such utility functions in his analysis of personal investment and consumption strategies. Another circumstance where more sophisticated levels of risk aversion may be appropriate concerns environmental resources. When there are a lot of these resources, one should behave with a relatively risk-neutral attitude, but as these resources become rarer, a more risk-averse attitude may be very appropriate to represent public values.

Practical Considerations for Selecting an Assessment Procedure

In principle, assessing the parameters of a value model (the single-attribute value or utility functions v_i or u_i , and the weights w_i or k_i) is only a matter of asking a

series of questions about indifferences and determining from this series of questions a set of n equations with n unknown and then to solve for the unknowns.

In practice, several procedures have emerged as being practical for assessing single-attribute utility and value functions and weights. We have found two types of assessment procedures especially useful: the equivalent-cost assessment procedure and the rating and swing weighting procedure.

Equivalent-Cost Procedure

This procedure is most appropriate, when the consequences can be described by simple numerical scales, when the single-attribute value functions are linear, and when the overall value model is additive. In that case, it is often possible to determine the unit equivalent cost of each attribute (e.g., the cost of an injury or of the loss of an acre of land) and to calculate the overall value of a vector of consequences as the sum of the equivalent costs of each attribute (calculated by the unit cost times the number of unit consequences for each alternative).

This procedure is similar to a benefit–cost model, except for two distinctions. First, the equivalent costs are obtained by asking decision makers for value trade-offs, instead of deriving them from market considerations and observed prices. Markets and prices can inform this dialogue, but other information is often used in this process as well. In many public decisions, there are no market prices for the consequences of some important objectives. Examples include lives lost, habitat destroyed, or jobs created. Second, several assumptions are typically checked to see whether the equivalent-cost model makes sense. These assumptions include the additivity and linearity previously mentioned and whether or not the cost attribute is a significant attribute in the value model. This latter point can be illustrated with a simple example.

In Figure 13.2, assume that the consequences of a decision can be displayed in terms of two attributes. One of these attributes is cost C . The other attribute Y measures an objective that one would like to minimize. For example, Y may be the acres of trees destroyed in building various facilities considered as alternatives for the problem. Suppose all of the consequences are as indicated in Figure 13.2a by dots. Also, suppose that one had assessed a complete value model as illustrated by the indifference curves with an equal value between adjacent curves. As seen from this figure, the range of the Y -consequences covers just the difference between adjacent indifference curves, whereas the range of the cost consequences covers about eight indifference curves illustrated on the C -axis. This implies that the cost implications in the decision, given the range of possible consequences, are about eight times more significant than the Y -consequences for that decision. Given this situation, it is quite reasonable to convert these Y -consequences into equivalent costs for the value model.

There are two ways one could obtain equivalent costs, illustrated respectively in Figures 13.2a and 13.2b. In Figure 13.2a, equivalent-cost values for all the Y -levels are determined and one “costs out” all of these down to the 0 level of the Y -attribute. If one did this for each of the point consequences in the figure (i.e., dots),

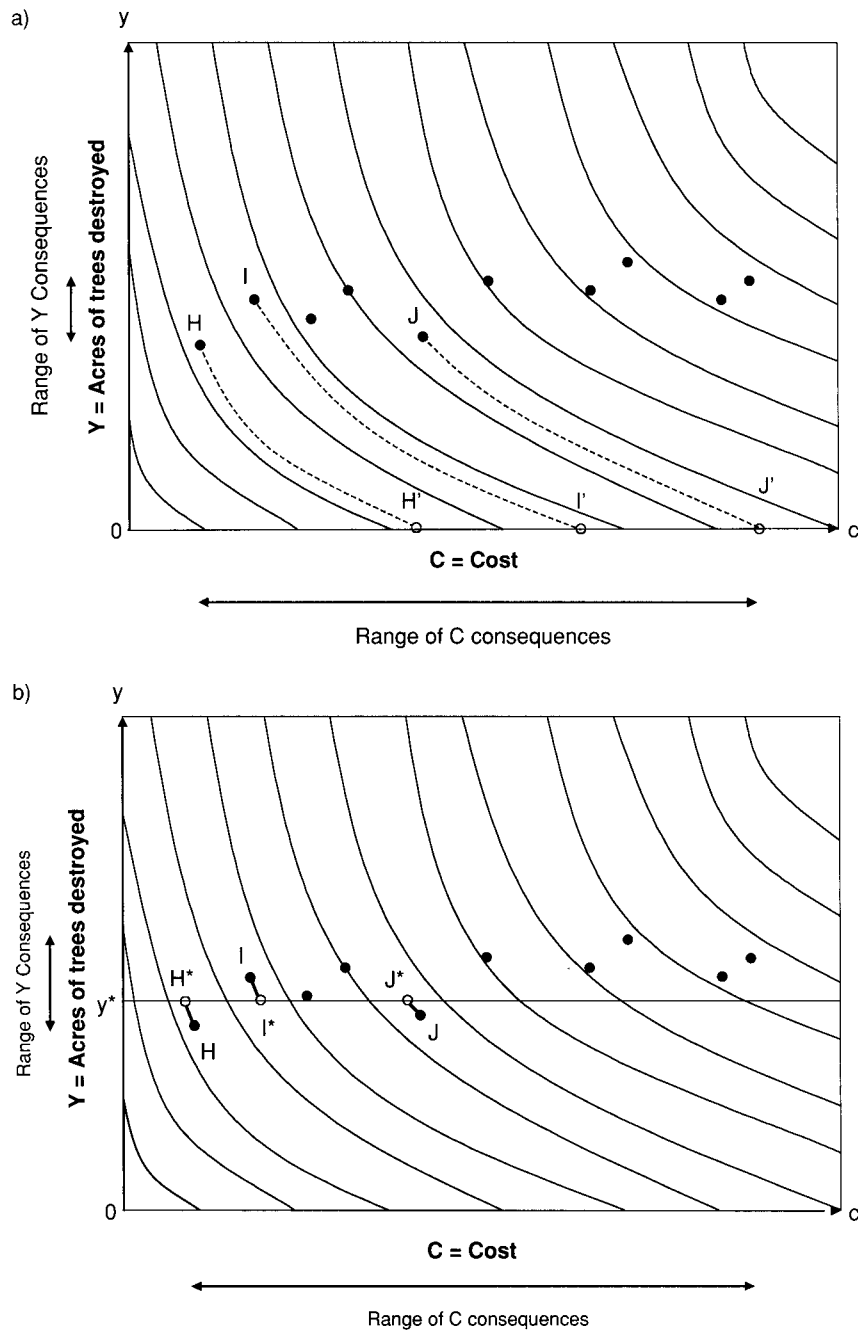


Figure 13.2. Converting two-attribute consequences to one attribute using a reasonable linear value tradeoff. The figure above, (a), represents nonlinear tradeoff over long range, and the figure below, (b), linear tradeoff over short range.

the corresponding equivalent consequence would be found by following the indifference curve down to the Y axis. Thus, for instance, the consequence labeled H would be converted to the consequence labeled consequences H' . Note that the attribute level of Y associated with H' is outside the range of the set of Y -levels for the original consequences.

Now consider Figure 13.2b. Here we draw a horizontal line through the consequences given Y set arbitrarily at a level y^* in the original range of the Y -levels. Now we can find the consequence equivalent to H , referred to as H^* , that is on the y^* line of consequences. Note that consequence H is much closer to consequence H^* than it is to consequence H' in Figure 13.2a. Thus, it is reasonable to use a linear function to translate consequence H into H^* . For instance, that transformation may state that an acre of trees has an equivalent value of D dollars. The same unit value tradeoff, meaning equivalent dollars per acre of trees, might be used for consequences I and J , illustrated in Figure 13.2b, to convert them to equivalently valued consequences I^* and J^* .

As a result of conversions such as this, an equivalent valued consequence with a y -level of y^* is found for each of the original consequences in the decision and it is naturally differentiated only on the cost attribute. Thus, only the equivalent costs of those consequences are necessary for differentiating the desirability of the alternatives. Also, as a result of the closeness of the equivalent consequences to the corresponding original consequences, the relative impact on the evaluation due to the assessed value tradeoffs to develop the cost-equivalent evaluations is much less than when all of the Y -levels are costed out to a level, where $y = 0$.

Now let us return to the general case where there are n objectives and attributes X_i other than cost C . Let a consequence be described by $(x_1, x_2, \dots, x_i, \dots, x_n, c)$, where c is the cost level. If all tradeoffs between X_i and C are linear (or approximately linear as in Figure 13b above), then an equivalent cost model is a practical value model and its construction is particularly simple. In this case, one needs to elicit only the equivalent costs for a unit consequence of each attribute to calculate the equivalent cost v as

$$v(x_1, x_2, \dots, x_i, \dots, x_n, c) = c + \sum w_i x_i, \quad (13.10)$$

where w_i is the equivalent cost of one unit of attribute X_i .

The assumptions necessary for an equivalent-cost model are often appropriate for major public policy decisions. We have used this simple value model in many recent decision analysis applications, including siting the first nuclear repository in the United States (Merkhofer and Keeney 1987), managing nuclear waste from powerplants (Keeney and von Winterfeldt 1994), retrofitting buildings to reduce earthquake risks (Benthien and von Winterfeldt 2001), managing potential health risks from electric powerlines (von Winterfeldt et al. 2004), and reducing risks of a dam failure (Keeney and von Winterfeldt 2005).

Rating and Weighting Procedure

In problems that involve softer attributes like “aesthetics” or attributes that do not have a cost equivalent, it is often practical to use a rating and weighting method, as described, for example in Edwards and Newman (1982) and von Winterfeldt and Edwards (1986). In this method the consequences of alternatives are rated on a scale (usually from zero to 100) reflecting their relative value on each attribute (100 being best). The attributes are then assigned weights, paying close attention to the relative ranges of the consequences in each attribute.

RATING CONSEQUENCES. Typically, this step starts by defining a low and a high consequence for each attribute and by assigning these two extremes a value of zero and 100. All other consequences are then rated between zero and 100 reflecting the degree of preference of each rated consequence relative to the lowest and highest ranked ones. When the attribute is defined on a continuous numerical scale, other methods such as bisection (dividing the scale into steps of equal changes of preference) can be used as well.

WEIGHTS. Weights can often be assessed for either units of the attributes or for ranges of the attributes. It is very important that weights are not assigned to objectives and attributes themselves as this can lead to significant biases (e.g., von Winterfeldt and Edwards 1986). If the attributes have linear value functions, it is easiest to assess the relative importance of a unit of each attribute. For example, in a recent study evaluating decision science programs at U.S. universities (Keeney, See, and von Winterfeldt 2006), we identified several attributes, including the number of refereed articles published in the past 5 years, the number of scholarly and popular books in prints, and the number of dissertations in the past 5 years. With linear attributes it is then easy to ask, for example: What is more important, one book or one refereed article and how much more important is it? Answers to these questions, provided in the study by eight scholars in decision analysis and behavioral research, proved to be quite consistent.

When attributes have nonlinear value functions, it is more appropriate to assign weights, referred to as “swing weights,” to the ranges of the attributes (von Winterfeldt and Edwards 1986). In this method, the decision makers are asked to consider a hypothetical alternative that is described by the worst consequences on all attributes. They are then asked which attribute they would like to change (“swing”) most from its worst level to its best level. Once this judgment is made, they are asked which attribute they would like to change next and so on until a complete rank order of these weights is established. With the rank order in place, one can assign an arbitrary weight of 1 to the lowest ranked attribute and ask the decision makers to scale the remaining weights based on the relative importance of those ranges. Alternatively, one can start with assigning an arbitrary weight of 100 to the highest ranked attribute and scale the remaining ones as fractions of 100.

When applying this method, it is important to cross check the results with some tradeoff questions. In particular, when cost is an attribute, it is important

to discuss the implications of these range weights for equivalent cost comparisons. For example, if the costs of decision alternatives range from \$200 million to \$300 million and with a possible loss of productive wetlands ranging from 0 acres to 100 acres, a decision maker may judge the latter range to be more important than the former. In that case it should be pointed out that this implies that it would be worth spending at least \$100 million dollars to avoid 100 acres of lost wetlands or \$1 million per acre. This may not be an unreasonable judgment, but it clearly depends on the type of wetlands, the existence of similar wetlands, and the number and rareness of species in this environment, etc.

Using Constructed Scales

With some objectives it is not possible to identify an attribute with a natural scale and, as a result, one needs to develop a constructed scale (see Chapter 7 in this volume). With a constructed scale, there are typically two to ten well-defined points on that scale. The basic task in a value model is to assign appropriate relative values to each of those levels. One needs to make sure that the values for these levels cover the range of possible consequences. For instance, if you have six possible levels of consequences defined on a constructed scale and wish to assign relative values from zero to one hundred to those consequences, is it not reasonable to have a value between zero and five assigned for five of those consequences and a value of one hundred assigned to the sixth. Unless there were no possible consequences between five and one hundred, it would be better if consequence levels on the constructed scale were chosen so that their utilities would correspond more closely to the amounts zero, twenty, forty, . . . , one hundred when consequences cover the full range of consequences.

Practical Considerations for Developing Value Models with Decision Makers

When the decision maker or decision makers are identified, their values should be incorporated in the value model. If the decision maker or decision makers have sufficient time and interest, then a complete value assessment could be done. However, sometimes they may have limited available time or they really do not understand why spending a significant amount of time expressing their value judgments is important. In these situations, it is still important to involve them at some level to obtain a few of the key value judgments. This both increases the decision makers' interest in the results of the analysis and their potential willingness to act on them.

Some decision makers consider it worthwhile to spend time expressing values for the analysis. In these cases, it is important to ask the most relevant value questions first. It is also reasonable to ask these value judgments from different perspectives that will facilitate consistency checks. We have found that if you can point out inconsistencies in values that decision makers consider important, they begin to understand why there is substantial value in clarifying their values for the decision. This certainly increases their willingness to participate.

For some decisions, especially those concerned with long-term consequences in governmental or other large organizations, it is not clear who makes the decision or perhaps even how the decision will be made. In these situations, the value model should represent values appropriate for the decision problem for the organization that will eventually make the decision. Input judgments about specific values can come from different individuals in the organization or constructed based on the knowledge of the organization by the analyst.

Generic Objectives

Obtaining an appropriate set of objectives for a decision problem requires effort and creativity. Indeed, Chapters 6 and 7 in this volume suggest procedures that might help people in this task. As a practical suggestion for specific types of problems, there are often generic sets of objectives that will apply to a wide range of decisions. For many public policy decisions, one wants to consider objectives in categories that might be referred to as costs, socioeconomic impacts, impacts on the natural environment, and health and safety. Sometimes this categorization is extended to include political consequences to or within the decision making organization and impacts on its relationship with public groups and other governmental organizations.

On any major decision, it is worthwhile to initially think of objectives from the viewpoint of various stakeholders concerned about a decision. For instance, if a large corporation is making a major decision, it is certainly worthwhile to consider objectives that deal with the impacts on shareholders, customers, employees, and management and then break these objectives into both short-term consequences and long-term consequences.

Constructing Value Models

In many situations, an analyst will begin to understand the values appropriate for the decision better than any one individual in the organization responsible for the decision. Part of this is because the analyst talks to many people throughout the organization and different individuals understand the appropriate values for different objectives, and part of it is because some individuals do not know how to express their values in a clear manner to be used in the value model. Analysts can help by combining values expressed throughout the organization and by improving the communication of values within the organization.

Decision makers will often indicate, in general terms, that they are risk averse with regards to consequences of a certain attribute. However, they do not seem to understand thoroughly the value questions concerning choices between lotteries, either because they do not understand well-defined probabilities or they cannot deal with the hypothetical circumstances that have a few consequences in a simple lottery when the real world has the possibility for numerous consequences. In these situations, it is reasonable for the analyst to translate the basic values of the decision maker in order to define parameters in the value model. This is

very consistent with the ideas of Payne and Bettman and their colleagues (Payne et al. 1992; Payne et al. 1999), regarding the constructive nature of building value models.

Focus on Key Values

There are many situations where the time available to interview the key decision makers about their values is limited, or when one is not sure how much time will be available. In both of these situations, it is critical to focus the interview process on the key value judgments that are presumed to strongly influence the results of that analysis. One of the most important aspects of this is to make sure that all the fundamental objectives of concern to the decision maker are included in the analysis. The other two types of critical value judgments typically concern the value tradeoffs between some key attributes and the levels of risk aversion appropriate for certain attributes. The intent is to quickly find out what is going on in the head of the decision maker so that the analyst can spend the appropriate time to construct a model.

Consider an analysis of different alternatives that address potential terrorist activities. Furthermore, assume that the potential loss of life is one of the main objectives. There are circumstances where the decision maker, when being asked about relative values for the life lost, might say that the first loss of life is worth at least a billion dollars, and the next few are worth more than \$100 million dollars a piece, but after 100 or so fatalities, the loss of additional life is worth about \$5 million in equivalent dollars. If these preferences are meant to reflect a utility function, it would be consistent with a strongly risk prone utility function. The important point here is to inquire about why these are the preferences of the person expressing them. It may be that the first few fatalities indicate failure of a major system, which will induce a large public response that results in major psychological and economic impacts. Once the terrorist acts have caused say 100 fatalities, then the indirect psychological and economic consequences will have been caused, and the additional impact of more fatalities is the additional loss of life. Understanding this as the reason for the preferences allows one to do a much better job of defining the objectives of the problem. Specifically, we recommend separating out objectives to indicate the indirect consequences of the terrorist act from the direct loss of life due to that act. Then the single-attribute value model over the fatalities should be linear if it concerns the loss of life and the impact on families and friends only. The psychological and economic consequences can then be defined more clearly and presumably values for different levels of those consequences can be better understood and assessed.

Bounding Value Parameters

For parameters that deal with value tradeoffs, such as the w_i in Eq. 13.1 and with risk aversion, such as the c in Eq. 13.9, it is often useful to assess a range that the decision maker considers appropriate to bound the parameter level. Then

one can easily perform a sensitivity analysis by varying the parameters within this range. In many situations, it turns out that the evaluation of the alternatives and the insights from the analysis do not depend on most of these parameters given their reasonable ranges. This allows one to focus on those parameters that are most important to the problem, which leads to the critical insights that can inform making better choices.

Understanding Values

One of the most practical and powerful of assessment tools is simply to ask the question “why?” Assessing a utility function or a measurable value function involves a series of questions. The intent is to understand what is in the mind of the person being assessed about values appropriate to the decision of concern. It is not hard to get an answer to many questions about values, but one wants to probe deeper to understand the fundamental value judgments that are leading to these responses. Inquiring about the reasoning for each response is very useful. The responses help both to build a better value model and provide a better analysis.

Suppose a decision is being made about placing natural gas pipelines through forested areas to bring natural gas to an urban area. One of the negative consequences might be the acres of forest destroyed in placing those pipelines. When asked a question about what is the value tradeoff between costs and an acre of forest destroyed, the respondents may say \$1 million is an equivalent cost to use for an acre destroyed. Asked why, the respondents may state that the price of the lumber for an acre of that forest is worth \$1 million. Alternatively, they may respond that the forest is the home to many animal species and the value of the habitat of those species is \$1 million. Or they may state that many individuals take a vacation in that area and enjoy the forest, and over the time that the forest is disrupted, the value of the loss of enjoyment by those individuals is \$1 million. Any of those responses suggest ways in which the values should be potentially separated out for the decision problem and/or ways in which the values are not complete. Indeed, if each of those three sets of implications were tied to the loss of an acre of forest, and if those stated value judgments were appropriate for evaluating the components of that loss, one might say that \$3 million was equivalent to the loss of an acre of that forest. The point is that by routinely asking the question “why,” one learns a lot more about the values.

Conclusions

When building value models an analyst has to answer several critical questions and, depending on the answers, has to make decisions about what type of value model to build and how to build it. The most important questions are:

1. Does the problem require a value function or a utility function? If a utility function is needed, should one assess a value function first?

2. Is the value function or the utility function additive? If not, what other aggregation model should be used?
3. Are the single-attribute value or utility functions linear and, if not, how should risk aversion be addressed?
4. Is an equivalent-cost model appropriate for the problem or should a rating and weighting method be used?

The answers to the questions and the implications for building a value model should be driven by the problem that needs to be solved and by the decision makers' needs and preferences. In our experience over the past two decades of building value models, we have found that, depending on the answers to these questions, two classes of models and some minor variants are often appropriate. The first class of value models is an equivalent-cost value function, possibly with a parametric exploration of risk aversion through an exponential transformation into a utility function. The second class of value models is a simple additive value function with unit or range weights and with no particular need to develop a utility function to account for risk attitudes.

The equivalent-cost model has proven valuable in many public policy applications, but it is also applicable in many decisions for private businesses, where the objectives are more limited (e.g., net profit, growth, and market share). It has been our experience that it is almost always useful to first build a value function and that a careful choice of fundamental objectives often justifies a linear or near-linear single-attribute value functions. If this is the case, and cost is a major part of the decision, then using linear value functions and unit costs to determine weights is both easiest to assess and to communicate. An example of this type of value model is used in the analysis of alternative policies to reduce electromagnetic field exposure from electric powerlines (see von Winterfeldt et al. 2004).

Several variants of this approach have been used in the past. For example, when constructed measures are used, it is better to assess weights using the equivalent values for the whole range of the attribute rather than for attribute units (see Keeney and von Winterfeldt 2005). When risk aversion is important, it is often useful to supplement the value function with a parametric analysis of a constantly risk averse utility function (see Merkhofer and Keeney 1987).

Using a simple additive value model without assessing equivalent costs is useful in "softer" problems when costs and uncertainties do not play an important part. Assessing the quality of consumer products is one example. Another example is a recent evaluation of decision science programs at U.S. universities (Keeney et al. 2006). These simple models are similar to the rating and weighting models first proposed by Edwards and his colleagues (see Edwards and Newman 1982; von Winterfeldt and Edwards 1986). However, it is important not to oversimplify the value model in these cases. In particular, we emphasize the importance of using fundamental objectives, natural or constructed measures, checks of additivity, and care in assessing unit or range weights.

As with the equivalent-cost model, it is useful to identify fundamental objectives and attributes that are additive and have linear or near-linear single-attribute

value functions. If some objectives appear nonadditive, we suggest redefining the objectives or adding objectives that can capture the reasons for the nonadditivity. Similarly, if some attributes have clearly nonlinear value functions, it is useful to probe the reasons for the nonlinearity and possibly redefine objectives or separate an objective into several parts. Regarding attributes, we encourage the use of natural measures to the extent possible. This helps address the factual part of the problem – the assessment of the degree to which alternatives satisfy objectives – and separates the value model from the factual model. When constructed measures are needed, it is important that these are carefully defined so that independent experts can arrive at similar judgments about how an alternative scores on the constructed measure.

We end with a caveat. This chapter is about using practical value models, meaning models that are good enough to gain all the important insights for a decision that might be expected or hoped for. We believe that often this does not require developing a state-of-the-art value model. It does, however, always require sound logic and judgments to justify the decision not to build a state-of-the-art model. Nothing in this paper is meant to suggest that a quick and dirty listing of poorly defined objectives and rating or weighting of these objectives is reasonable or justifiable.

There are times, especially for very complex and important decisions, when one should develop a state-of-the-art value model. This requires explicit assessments to justify all assumptions, to select the set of fundamental objectives, to choose attributes, to select all utility or value functions (both in single-attribute and multiattribute form), to assess the parameters, and to perform consistency checks for all elements of a value assessment.

REFERENCES

- Barron, F. H., von Winterfeldt, D., and Fischer, G. W. (1984). Empirical and theoretical relationships between value and utility functions. *Acta Psychologica*, 56, 233–244.
- Bell, D., and Raiffa, H. (1988). Marginal value and intrinsic risk aversion. (1988). In D. Bell, H. Raiffa, and A. Tversky (eds.) *Decision making: Descriptive, normative, and prescriptive interactions*. Cambridge, UK: Cambridge University Press, 384–397.
- Benthien, M. and von Winterfeldt, D. (2001). Using decision analysis to improve seismic rehabilitation decisions. Working Paper WP-01–01, Institute for Civic Enterprise, University of Southern California, Los Angeles, CA.
- Dyer, J. S. and Sarin, R. (1979). Measurable multiattribute value functions. *Operations Research*, 22, 810–822.
- Edwards, W., and Newman, J. R. (1982). *Multiattribute evaluation*. Beverly Hills, CA: Sage.
- Fishburn, P. C. (1965). Independence in utility theory with whole product sets. *Operations Research*, 13, 28–45.
- Fishburn, P. C. (1970). *Utility theory for decision making*. New York: Wiley.
- Hammond, J. S., Keeney, R. L., and Raiffa, H. (1999). *Smart choices: A practical guide to making better decisions*. Boston, MA: Harvard Business School Press.
- Keeney, R. L. (1981). Analysis of preference dependencies among objectives. *Operations Research*, 29, 1105–1120.
- Keeney, R. L. (1988). Building models of values. *European Journal of Operational Research*, 37, 149–157.

- Keeney, R. L. and Raiffa, H. (1976). *Decisions with multiple objectives: Preferences and value tradeoffs*. New York: Wiley. Reprinted in 1993 by Cambridge University Press.
- Keeney, R. L., and von Winterfeldt, D. (1994). Managing nuclear waste from power plants. *Risk Analysis*, 14, 107–130.
- Keeney, R. L., See, K., and von Winterfeldt, D. (2006). Evaluating interdisciplinary academic programs: With applications to U.S. graduate decision programs. *Operations Research*, 25, 1–16.
- Keeney, R. L., and von Winterfeldt, D. (2005). *Evaluation of Capilano-Seymour tunnel project alternatives*. Technical Report, submitted to the Greater Vancouver Regional District, Vancouver, Canada.
- Krantz, D. H. (1964). Conjoint measurement: The Luce-Tukey axiomatization and some extensions. *Journal of Mathematical Psychology*, 1, 248–277.
- Krantz, D. H., Luce, R. D., Suppes, P., and Tversky, A. (1971). *Foundations of measurement, Volume I*. New York: Academic Press.
- Luce, R. D. and Tukey, J. W. (1964). Simultaneous conjoint measurement: A new type of fundamental measurement. *Journal of Mathematical Psychology*, 1, 1–27.
- Merkhofer, M. L. and Keeney, R. L. (1987). A multiattribute utility analysis of alternative sites for the disposal of nuclear waste. *Risk Analysis*, 7, 173–194
- Meyer, R. F. (1970). On the relationship among the utility of assets, the utility of consumption, and investment strategy in an uncertain, but time invariant world. In J. Lawrence (Ed.). *OR 69: Proceedings of the Fifth International Conference on Operational Research*. London: Tavistock Publications.
- Payne, J.W., Bettman, J.R., Coupey, E., and Johnson, E.J. (1992). A constructive process view of decision making: Multiple strategies in judgment and choice. *Acta Psychologica*, 80, 107–141.
- Payne, J. W., Bettman, J. R., and Schkade, D. A. (1999). Measuring constructed preferences: Towards a building code. *Journal of Risk and Uncertainty*, 19, 243–270.
- Pratt, J. W. (1964). Risk aversion in the small and the large. *Econometrica*, 32, 122–136.
- von Neuman, J., and Morgenstern, O. (1947). *Theory of games and economic behavior* (2nd ed.). Princeton, NJ: Princeton University Press.
- von Winterfeldt, D. (1979). *Functional relationships between risky and riskless multiattribute utility functions*. Technical Report No. 79–3. Social Science Research Institute, Los Angeles: University of Southern California.
- von Winterfeldt, D. and Edwards, W. (1986). *Decision analysis and behavioral research*. Cambridge, UK: Cambridge University Press.
- von Winterfeldt, D., Eppel, T., Adams, J., Neutra, R., DelPizzo, V. (2004). Managing potential health risks from electric powerlines: A decision analysis caught in controversy. *Risk Analysis*, 24, 1487–1502.