Modeling Values for Counter-Terrorism Analysis

By Ralph L. Keeney

Fuqua School of Business

Duke University

October 28, 2005

Abstract

Decisions are made to achieve objectives. A value model unambiguously represents objectives in a quantitative manner. Hence, a value model both guides thinking and provides a basis for analyzing alternatives to best meet the desired objectives. This paper proposes the development of value models to support counter-terrorism efforts. It summarizes the theory and procedures to develop value models for the Department of Homeland Security and for terrorist organizations. The later is useful to both design counter-terrorism alternatives and suggest possible terrorist priorities and actions. An example that develops a terrorist value model for the theft and misuse of plutonium is presented. Several uses of value models for counter-terrorist activities are discussed and suggestions for developing such value models are outlined.

1. Introduction

The US Department of Homeland Security (DHS) was created as one response to the terrorist events of September 2001. The DHS annual budget is approximately $40 billion. A large portion of these funds are allocated to implement decisions to make our country and its citizens safer from terrorism. These decisions to counter terrorist activities are obviously important and complex. The complexity results from the multiple objectives that one wishes to achieve, major uncertainties about how well different alternatives will achieve them, the wide range of potential alternatives that may be appropriate, and a dedicated adversary (i.e. terrorists) that can take actions to alter the intended consequences. One cannot systematically and logically understand and integrate all relevant information informally in one's head and make an informed choice using intuition alone. Insightful analysis can be of great help, but to be useful to inform those decisions, any analysis must address the complexities explicitly.

The best choice in any decision situation depends on three critical factors: the alternatives that are identified, the objectives and their priorities for what one wishes to achieve, and the likelihoods describing how well the different alternatives meet each of the objectives. This paper focuses on the less considered critical factor, namely the objectives and their priorities for what the Department of Homeland Security should strive to achieve.

In the more than four years that have passed since September 2001, we should have produced a very clear statement of our nation's objectives for counter terrorism activities. In my knowledge we have not. We have all of the information to produce such a statement and logically prioritize the objectives. That information is based in our nation's values summed up by life, liberty, and the pursuit of happiness. The task is to define operationally what these values imply in terms of objectives relevant to counter terrorism activities. Relative to the DHS budget, the cost of doing this well would be miniscule. We should do this now.

The objectives and their priorities define what we hope to achieve, so they should guide all of our decisions made to counter terrorism. If we don't have a clearly understood statement of what we hope to achieve, how can we hope to achieve it? If we do have such a statement, we should be able to much better achieve it.

An analysis of alternatives to counter possible terrorist attacks has certain features in common with an analysis of possible natural disasters such as floods, hurricanes, and earthquakes. Both types of decisions have a myriad of objectives such as preventing loss of lives, limiting the disruption of lives, and minimizing economic costs. There are also major uncertainties about possible consequences in both situations. We cannot forecast the exact locations or magnitudes of natural disasters, nor can we describe their consequences precisely. These uncertainties are due to a lack of a complete understanding of Mother Nature and of how people will react to a natural disaster. With possible terrorist acts, the major uncertainties are due to a lack of knowledge about possible terrorist acts concerning when, where, and how they will be carried out and about how people will react to these acts. Unlike nature, terrorists can adapt their actions

to account for decisions that we have taken to protect our citizens and country. In other words, we have a thinking adversary who attempts to render our actions ineffective.

The obvious competition between DHS and terrorists suggests the possible usefulness of game theory as a foundation for analysis (Luce & Raiffa, 1957). I think the spirit of incorporating the opponent's actions in appraising our best alternatives is necessary for good analysis. However it is not reasonable to assume that we have common information, or more importantly, that we can consider numerous possible actions and reactions into the future. A decision analysis along the lines Figure 1 that accounts for an initial DHS decision, a terrorist response, then a second DHS decision, a second terrorist response, and then the uncertainty about public reaction may often be sufficient for useful analysis.

For such an analysis, or indeed any analysis of DHS actions, a clear understanding of the objectives to be achieved is important. The theory and procedures to build value models that state objectives and quantify priorities for DHS to use in such analyses are well developed. Unfortunately, this body of knowledge is not routinely used, and the state of practice in developing such models in DHS is apparently woefully inadequate.

To gain insight about possible terrorist actions in such models, it is useful to have a value model representing terrorist interests and priorities. The theory is the same as that for DHS, but the procedures to construct such a model need to be altered, as terrorists would not likely willingly cooperate with analysts building such a value model.

This article is about building and using value models for the analysis of terrorist activities and counter-terrorist decisions. The basis for this article rests on four facts: (1) decisions should depend on what the decision makers hope to achieve, (2) value models make these objectives explicit and quantify what those decision makers hope to achieve, (3) we have the theory and procedures to construct quality value models, and (4) knowledge of such values is critical for an informed choice. Section 2 presents a summary of the theory for constructing models of values. Section 3 describes procedures for implementing this theory to develop specific value models for DHS and for terrorists. Section 4 uses this theory and procedures to develop a value model concerning the potential theft and misuse

of plutonium by terrorists. Section 5 describes several uses of value models for counter terrorist activities. Section 6 presents a summary and recommendations.

2. Theory For Models of Values

Another term for a value model is an objective function. The term value model has the advantage that it accurately suggests that the value model is constructed using the same processes necessary to construct any type of analytical model. The analyst decides on the variables to use, verifies appropriate relationships among those variables to get a mathematical representation of interest (a quantification of preferences in this case), and then specifies some parameters for the model using information that is available or can be gathered.

There are several published sources that present the mathematical foundations for constructing value models (Krantz et al, 1971; Keeney and Raiffa, 1976; and von Winterfeldt and Edwards, 1986). Because of the significant uncertainties in terrorist and counter-terrorist decisions, the appropriate type of value model is a utility function developed consistent with utility theory (von Neumann and Morgenstern, 1947). A utility function has the property that the expected utility of any alternative indicates its desirability. Specifically, the utility function is constructed such that alternatives with higher expected utilities are preferred to alternatives with lower expected utilities.

Details about the construction of value models are discussed elsewhere (Keeney, 1988). The process can be effectively summarized with five steps;

1. Specifying objectives,
2. Identifying attributes for those objectives,
3. Relating and integrating objectives,
4. Specifying relative preferences for different levels of the single attributes,
5. Defining the value tradeoffs that prioritize the different objectives

Specifying Objectives. The first step in developing a value model is to get a comprehensive list of the fundamental objectives for the decision being considered. At the highest level of an organization, such as DHS, there should be a common set of

strategic objectives that should provide guidance for all decisions of the organization. Specific decisions may focus on more specific objectives that help achieve those strategic objectives, so specific decisions may require a set of fundamental objectives designed for that decision. As an example, one strategic objective of the DHS would be to minimize loss of life to terrorism. One of the ways this can be done is to reduce possible terrorist acts, which is a means objective for the strategic objective of preventing loss of life due to terrorism. However, in a specific decision context focused on reducing terrorist acts, one fundamental objective may be to minimize terrorist acts.

There are five desirable properties for a good set of fundamental objectives. Summarizing from Keeney (2006) they are the following:

- Complete: All of the important consequences of alternatives in the decision context can be adequately described in terms of the set of fundamental objectives.
- Non-redundance: The fundamental objectives should not include overlapping concerns.
- Concise: The objectives and subobjectives, that more precisely define the objectives, should be the minimum number appropriate for quality analysis.
- Specific: Each objective should be specific enough so that the consequences of concern are clear and useful attributes can be readily selected or defined.
- Understandable: Any interested individual knows what is meant by the objectives.

Attributes to Measure Objectives. For each of the fundamental objectives, or its subobjectives, a measure referred to as an attribute is needed to indicate the degree to which that objective is met. For the example above concerning minimizing the loss of life to terrorism, an attribute would be the number of individuals dying due to terrorism annually in the US. For the objective of minimizing property damage due to terrorism, an attribute might be “economic value of the property lost in billions of dollars”. In cases like these, the attributes seem obvious. In other cases they are not so obvious, even though the objective is very important. One objective related to terrorist and counter terrorist activities might be to minimize the fear and despair that it causes. Another would be to minimize the restrictions on civil liberties. For objectives such as these, clear and hard thinking is necessary to develop an attribute, or perhaps a set of attributes for

subobjectives, that can adequately describe the respective consequences of any alternative (see Keeney and Gregory, 2005).

General Structure of the Value Model. Once a set of fundamental objectives and associate attributes are defined, the next step is to select the general structure of an appropriate value model. To describe this process, it is useful to introduce some notation. Let us label the identified fundamental objectives as O_i, i=1, ..., n. and the associated attributes X_i, i=1, ..., n. With this notation, a possible consequence of any chosen alternative can be written as $(x_1, ..., x_n)$ or more compactly as **x** where x_i is a specific level of X_i. We can then use the notation u(**x**) to indicate the utility function which is the value model of interest.

The value model combines preferences for the various levels of the n attributes in a function that reflects the values of the decision maker or the decision maker's organization. This model is analogous to building a factual model of, for instance, the response of emergency medical services to terrorist acts. In that situation the model would include the location of ambulances at the time of the attack, possible consequences of the attack and the needs for service, and the ability of medical facilities to handle casualties. To combine the attributes in the value model, one uses values independence concepts and preference relationships analogous to probabilistic independence and logical relationships used in factual models. The main types of value independence concepts used are utility independence, preferential independence, and additive independence that are described in detail in Keeney and Raiffa (1976).

If the set of fundamental objectives satisfy the desirable properties discussed above, very likely either the additive or multiplicative utility function is appropriate. These are respectively

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

and

$$1+ ku(x_1,...,x_n) = \prod_{i=1}^{n} \{(1+kk_i u_i(x_i)\}, \; k \neq 0, \qquad (2)$$

where u is the overall value model, the u_i are single attribute utility functions, and the k_i are scaling constants. The additive utility function is a limiting special case of the multiplicative utility function (2), since the multiplicative utility function approaches the additive utility function as the constant k in (2) approaches 0. The original work that developed the additive utility function is Fishburn (1965).

Specifying Single-Attribute Utility Functions. Single-attribute utility functions are constructed in a manner analogous to constructing multiattribute utility functions. In this case, the common concepts used to specify the form of single-attribute utility functions are risk aversion, risk neutrality, and risk proneness (Pratt, 1964). A specific case is when constant risk aversion holds which implies the simple exponential form

$$u_i(x_i) = a_i + b_i(-e^{-c_i x_i}), b_i > 0, c_i > 0, \quad (3)$$

where a_i, b_i, and c_i, are constants to be assessed.

Value Tradeoffs. Values conflict when there are circumstances where we can enhance the achievement of one objective only if we accept a degradation in the achievement of another objective. The judgments made about how much we are willing to give up in terms of achieving one objective in order to better achieve a second objective by a specific amount is referred to as value tradeoff. Value tradeoffs are at the heart of making informed choices. They are obviously both difficult and important to make, but neglecting them does not result in better choices.

A situation we have often read about recently concerns the security procedures at our nation's airports. There are concerns that any tightening of the procedures restricts civil liberties of the flying public. At the same time, such a tightening should result in the smaller likelihood of a successful terrorist act and hence might reduce the subsequent loss of lives. Value tradeoffs indicate whether any particular restriction of civil liberties is worth the potential benefits in safety. A small restriction of civil liberties may easily be worth a large improvement in safety, and a significant restriction of civil liberties may not be justified by a small improvement in safety. Hence, the specification of value

tradeoffs does not involve choosing one value, or its associate objective, over another, but rather it is the degree of achieving each objective that matters. These value tradeoffs were studied in a survey by Viscusi and Zeckhauser (2003).

To determine appropriate value tradeoffs requires finding even swaps, meaning pairs of different consequences that we would be willing to evenly trade for each other. We might conclude, as above, a willingness to swap the elimination of terrorist risk causing one expected fatality for an elimination of attorney-client privilege for 500 individuals in federal custody because of terrorism concerns. To make such a value tradeoff requires thinking about who and how many people are affected by each consequence and how severe each of the consequences are. Individuals with different perspectives should participate in discussions to identify an appropriate value tradeoff or a range of acceptable value tradeoffs. The reasons supporting different viewpoints should be clearly expressed so others can understand and appraise them.

3. Assessing the Value Judgments to Specify Value Models.

This section concerns how to assess the value judgments necessary to implement the five steps outlined in section 2. Although interaction between the steps is beneficial in practice, it is useful to present the ideas sequentially. We separately discuss the cases of value models for DHS and for terrorists.

Assessing Value Models for DHS. The basic way to identify objectives is to ask individuals what objectives are appropriate. Listing a few objectives is of course simple, but listing a quality set of objectives appropriate for DHS problems is difficult. A set of fundamental objectives that meets the criteria discussed in section 2 requires hard thinking, contributions by many individuals, and a subsequent in-depth appraisal and improvement.

To get a comprehensive list of objectives, many individuals concerned about terrorism should be asked for a list of such objectives for our nation. These individuals should include DHS officials, members of organizations with responsibilities related to terrorism, elected officials, and the public. One needs to stimulate the thought processes

of each individual and then take the union of all of their listed objectives. There are several devices suggested in Keeney (1992) to stimulate the creation of objectives.

After a broad set of objectives is developed, it is then necessary to organize them. This includes identifying the set of fundamental objectives in the decision context. It also involves identifying means-ends relationships, as following these relationships will help identify the fundamental objectives. For example, one objective of a screening process in an airport is to ensure a thorough search of each potential terrorist. This is a means to the objective of maximizing the likelihood that any potential terrorist is identified during screening. This is a means to reduce the likelihood that the potential terrorist will have access to the airport lounges and aircraft, which is a means of reducing the likelihood of a terrorist act at those locations. These are naturally means to minimizing the loss of life, injuries, and property damage. Each of these latter objectives is fundamental in this context, but they also contribute to subsequent fear that the public feels, which is also a fundamental objective in such situations. A preliminary list of objectives and attributes for DHS is given in Table 1 to indicate, by example, what we should create.

The development of attributes to measure each of the fundamental objectives is a more technical task than identifying and organizing objectives. Analysts knowledgeable about the intent of each of the objectives can usually develop useful attributes for each of the objectives. Guidelines for this process are found in Keeney and Gregory (2005).

If the criteria of a good set of fundamental objectives listed in section 2 are met, an appropriate attribute utility function should usually be the additive utility function (1). Details on why the additive utility function should be appropriate and procedures to develop other utility functions are discussed in Keeney (1992).

The single-attribute utility functions should have a shape that represents the preferences for the decision problem. It is usually appropriate to choose a functional form for that shape, such as the exponential form in (3). Then the scaling constants a_i and b_i can be set to scale the function from 0 to 1, which is appropriate scaling for the additive or multiattribute utility functions. The additional constant c_i in (3) can be determined by

finding two simple, hypothetical alternatives that are indifferent and differ only in terms of the attribute being considered. One of these alternatives must involve uncertainties, represented by specifying the probabilities of different consequences. Then, equating expected utilities using (3), one equation with the one unknown c_i is created. This equation can be solved to yield c_i.

Value tradeoffs, represented in the utility functions (1) and (2) by the k_i. are determined in a similar manner by finding pairs of consequences that are indifferent. Then, one can equate the utilities of those indifferent consequences using the chosen utility function. With n pairs of such consequences, one can create a system of n equations with n independent unknowns, namely the k_is in (1) or (2). In the case of the multiplicative utility function, the constant k can be calculated from the set of k_is. In the case of the additive utility function, there are only n-1 independent constants, so only n-1 pairs of indifferent consequences are necessary.

Assessing Value Models for Terrorists. The information that one needs to obtain a utility function for terrorists is exactly the same as for DHS. However, it is necessary to alter the procedure for obtaining the information as we cannot directly assess values from terrorists. It is worthwhile to point out that this circumstance is not unique: we are often interested in the values of parties to whom we cannot usually directly speak about their values. These include situations involving competitors in the business environment, preferences of customers, and negotiating with other parties. Certainly the objectives of these other parties are critical to both guide the creation of alternatives and to evaluate our own choices (Raiffa, 1982).

The objectives for a terrorist organization overall or for a specific terrorist activity can be elicited from numerous individuals. These would include government officials, scientists and analysts who study terrorism, members of the intelligence community who investigate terrorist organizations and their activities, and expatriates from nations known to harbor terrorists. In some cases, we can obtain objectives from the terrorists themselves, such as when terrorists who are in jail express their concerns and reasons for

being a terrorist. Terrorists who are not in jail sometimes articulate their concerns in the media through interviews or letters.

Although some of the objectives of terrorists might be exactly the opposite of the objectives of Homeland Security listed in Table 1, there would likely be other objectives that are different. These might include what terrorists may refer to as a better balance of support for both the Palestinian and Israeli sides on their issues, more respect for the Muslim religion, and less interference by Western nations in the internal affairs of other nations.

The specification of attributes for each of the fundamental objectives of terrorists would essentially follow the same procedures as for Homeland Security. Once the objectives are specified, selecting appropriate attributes requires hard and serious thinking, but it does not involve the level of creativity necessary to specify objectives.

Selecting a functional form for the utility function for a terrorist organization or terrorist decision is basically a choice that follows partly from the quality of a set of objectives. If the fundamental objectives meet the criteria laid out in section 2, then an additive utility function should be appropriate. .However, given the uncertainties we would naturally have about the preferences of the terrorists, a sensitivity analysis using the multiplicative utility function would be useful to investigate how potential actions of terrorists may differ if their preferences are different from our understanding of them represented by the additive form.

Value judgments will be necessary to specify the scaling constants in the single attribute utility functions and the value tradeoffs to characterize the scaling constants in the multiattribute utility functions. For these judgments, we need to rely on any of the individuals or groups who have knowledge of the preferences of the terrorist organizations. In using such a value model, it would be appropriate to do a sensitivity analysis that varies these scaling constants to examine possible implications of terrorist viewpoints different from those that we perceive.

It is important to acknowledge that value models are designed to be prescriptive, meaning to facilitate making better decisions. The manner in which it would be used in the Department of Homeland Security is prescriptive. However, from the viewpoint of the DHS, a model of terrorist values would be used to help describe how they may act. The shortcomings of utility theory for describing decision behavior are well known (Kahneman & Tversky, 1979; Kahneman, Slovic, and Tversky, 1982; von Winterfeldt and Edwards, 1986). On the other hand, when the stakes of the decision are as important as they are for the terrorists in such situations, it is reasonable to assume that they will develop and carefully choose those alternatives that have the best chances to achieve their objectives. Thus, a value model that characterizes what we think they wish to achieve, should facilitate our thinking and analysis of possible terrorist actions.

4. An Illustrative Value Model of Terrorist Preferences

This section illustrates the approach for developing a value model of terrorist preferences to help describe their possible actions. The problem context concerns the possible theft of nuclear material. A set of terrorist objectives and associated attributes are first discussed. The assessment of a utility function over these attributes is then illustrated. Potential uses of this utility function are discussed in Section 5. All judgments about objectives, attributes, and the utility function were assessed from knowledgeable scientists familiar with terrorist activities in the mid-1970s (Keeney, 1977).

The Terrorist's Objectives and Attributes. Our concern is an activity involving a single terrorist whose main objective is to obtain plutonium to create a bomb or aerosol device. A class of alternatives for doing this involve stealing nuclear material in a liquid solution. To help select among the alternatives, we assume that the terrorist has three objectives:

(1) maximize the amount of plutonium obtained,

(2) maximize the purity of the material stolen,

(3) minimize the personal radiation danger.

Reasons for the first and third objectives are obvious. To obtain plutonium, stolen material must be processed to remove impurities. Hence, objective two accounts for the time and effort needed to process the stolen material.

To measure the degree to which these objectives are achieved, our experts selected respectively the following three attributes:

X_1 = plutonium extracted, measured in grams (g.Pu.),

X_2 = purity of stolen material, measured in grams of uranium per liter of material (g.U./L.),

X_3 = radiation dose, measured in Rads./hr.

There is no claim that these three objectives capture all of the terrorist's concerns. For instance, the weight and volume of the material stolen, which influences the ease of the theft, and the length of time of exposure are not explicitly considered. Such considerations are left out for simplicity.

Before specifying a utility function, we need to determine ranges for the attributes. The ranges in Table 2 were chosen because they cover a wide range of potential targets for thefts. The designation of worst and best are taken from the terrorist's viewpoint.

General Structure of the Terrorist's Utility Function. Now we illustrate the assessment of a utility function $u(x_1, x_2, x_3)$, where x_i is a specific level of attribute X_i, i=1, 2, 3. Thus, for instance, u(200, 30, 10) is the utility to the terrorist of stealing 200 grams of plutonium with purity 30 g.U./L. which emitted 10 Rads./hr. The utility function can rank for this adversary all possible targets characterized by the three attributes.

To obtain the general structure for the terrorist's utility function, one investigates the appropriateness of specific assumptions about the preference structure. For instance, one ascertains whether value tradeoffs between any two attributes depend on the level of the third. If not, it is said that preferential independence holds. This assumption did hold. In addition, it was determined that X_1 was utility independent of the other two attributes,

meaning that the relative preferences for the amount of plutonium stolen did not depend on the other two attribute levels. These are the assumptions necessary to imply that either an additive utility function (1) or the multiplicative utility function (2) is appropriate for our illustrative example. The specific choice of which one to use depends on the answers to assessment questions.

The Single-Attribute Utility Functions. To assess the utility function for the amount of plutonium obtained from the theft, we fixed the other two attributes at convenient levels. The utility function u_1 was scaled from 0 to 1 by

$$u_1(10) = 0 \qquad (4)$$

and

$$u_1(2500) = 1, \qquad (5)$$

since 10 g. and 2500 g. are the worst and best, respectively, from the terrorist's viewpoint. We asked a knowledgeable scientist what amount for certain is indifferent to a lottery yielding a one-half chance at 2500 g. and a one-half chance at 10 g. First considering 1500 g. for certain or the lottery, the 1500 g. was preferred. Similarly, 1000 g. for sure was preferred to the lottery. However, 500 g. was not as desirable as the lottery. Finally, 800 g. was chosen as indifferent. For the terrorist being considered in this illustration, this means that 800 g. of plutonium for sure is indifferent to a fifty-fifty chance at either 10 g. or 2500 g. The implication of this is that the utility $u_1(800)$ assigned to 800 g. must equal expected utility of the lottery so

$$u_1(800) = 0.5\, u_1(10) + 0.5\, u_1(2500) = 0.5. \qquad (6)$$

Using (4) to (6), we have the three points on the scale in Figure 2. Using these, we could assess additional points on the graph and then fit a curve. Alternately, we can investigate the perceived terrorist's attitude toward risk and then fit a curve to the three points (see

Pratt, 1964). For this illustration, we fit the constantly risk averse exponential utility function

$$u_1(x_1) = 1.26(1 - e^{-1.58(\frac{x_1}{2500})}) \tag{7}$$

using (4) to (6). Similar procedures were used to assess u_2 and u_3 shown in Figure 2. We have assumed here that with attributes X_1 and X_3 fixed, an impurity of 10 g.U./L. for sure is indifferent to a lottery with a one-half chance at a solution with 0g.U/L. (i.e. completely pure) and a one-half chance at a solution with 333g.U./L. impurity. The reason is that the time and effort to purify a solution with 10 g.U./L. may be about half that to purify a solution of 333 g.U./L. Since impurity of more than 50 g.U./L. requires essentially the same effort to purify, we have assumed constant utility for impurity greater than 50 g.U./L. The resulting utility function is

$$\begin{aligned} &u_2(x_2) = 0.039(e^{0.066(50-x_2)} - 1), x_2 \le 50, \\ &u_2(x_2) = 0, x_2 > 50. \end{aligned} \tag{8}$$

For attribute X_3, we found that a radiation dose of 60 Rads./hr. is equivalent to a lottery yielding an equal chance at either a dose of 100 Rads./hr. or 0 Rads./hr. The utility function is

$$u_3(x_3) = 1.78(1 - e^{+0.822(\frac{x_3-100}{100})}). \tag{9}$$

Value Tradeoffs. To determine the value tradeoffs, which allow us to evaluate the k_i's in (1) or (2), it is necessary to consider two attributes at a time. To illustrate, let us suppose that the impurity level is fixed at x_2 in the top of Figure 3 and consider consequences of the form (x_1,x_2,x_3) for this fixed level of x_2. We ask which is preferred between (1000, x_2, 100) and x_1=(10, x_2, 0). If the latter is preferred, then u(10, x_2, 0) must be greater than u(1000, x_2, 100). By moving either x_1 or x_3, we search for a point of indifference. Suppose (2500, x_2, 100) is indifferent to (10, x_2, 0), then

$$u(2500, x_2, 100) = u(10, x_2, 0). \tag{10}$$

Using the scaling convention in (7) to (9), we can set x_2=333 and evaluate (10) using either (1) or (2) to find

$k_1 = k_3$. (11)

Considering value tradeoffs between X_2 and X_3 with plutonium fixed at x_1, suppose we find that (x_1, 0, 100) indifferent to (x_1, 333, 80) as shown at the bottom of Figure 3. We can then set x_1=10 and equate the utilities of these two consequences using either (1) or (2) to yield

$k_2 = k_3 u_3(80)$. (12)

Equating $u_3(80)$ from (9), we find (12) becomes

$k_2 = 0.27k_3$. (13)

In order to select a third (and perhaps fourth) equation to be used with (11) and (13) in evaluating k_1, k_2, and k_3 (and perhaps k), we consider the choice between the two lotteries

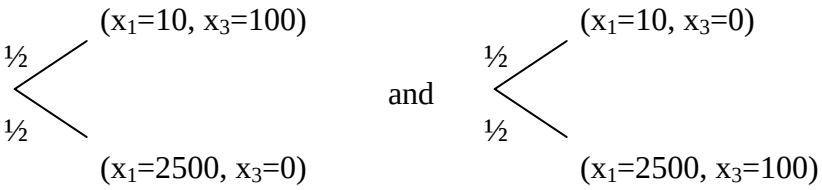

where x_2 is assumed to be fixed. Indifference implies the additive form and then $k_1+k_2+k_3=1$, which can be combined with (11) and (13) to solve for the k_i's.

However, in this case, the second lottery was preferred using the reasoning that the terrorist certainly wants the rewards for being exposed to all that radiation. This implies the multiplicative utility function (2) is appropriate. Evaluating (2) at the best possible consequence (x_1=2500, x_2=0, x_3=0), we find

$k+1=(kk_1+1)(kk_2+1)(kk_3+1)$. (14)

Finally, by varying the probability p in the lottery

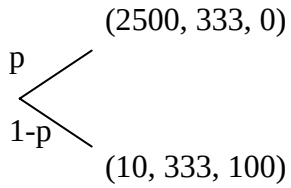

and comparing it to the consequence (2500, 333, 100), it was assumed that the two are indifferent for the terrorist when p=0.7. Then, equating their utilities using (2), we find

$$k_1 = 0.7(k_1+k_3+kk_1k_3). \quad (15)$$

Solving (11), (13), (14), and (15), we find

$$k_1 = 0.68,\ k_2 = 0.18,\ k_3 = 0.68, \text{ and } k = -0.842. \quad (16)$$

The terrorist's utility function u for the values postulated is given by (2) where (9) to (11) define the u_i's and (16) summarizes the scaling factors. This utility function is scaled from 0 to 1 by

$$u(10, 333, 100) = 0 \quad (17)$$

and

$$u(2500, 0, 0) = 1. \quad (18)$$

5. Uses of Value Models.

There are many important uses of both the process of developing a value model and of the resulting value model that is developed. The process enhances communication about the values that we consider important, a convergence to those that we want to use in evaluating our decisions, and an understanding of what is important and why. Once the value model for DHS is constructed, it can be used to facilitate the creation of alternatives and to evaluate DHS choices. Once a value model for terrorists is constructed, it can be used to provide insight into the actions that terrorists maybe contemplating.

This section suggests several ways which value models would be useful for DHS. The first use, understanding our values, is adapted from Keeney (2001) and discussed in much

more detail than the others, both because understanding our values is necessary for any of the other uses and the other uses are more generally understood by analysts and individuals with responsibilities for counter-terrorism decisions.

Understanding Our Values. Is our right to life, liberty, and the pursuit of happiness important? Of course. Do we want to maintain all of our civil liberties? Sure. Is freedom of speech essential to our society? Definitely. Can we simultaneously maximize each of them? No. These values, that provide some of the foundation of our nation, are sometimes in conflict.

Values conflict when we, meaning our nation, can enhance the manifestation of one value only if we accept a degradation in the manifestation of another value. Because values conflict, decisions concerning them involve value tradeoffs. Such decisions do not involve choosing one value over another value, rather it is the degree of each that matters.

We all care about the consequences of decisions that affect our basic values, so there is merit in trying to make the necessary value tradeoffs explicitly in an informed manner. Yet, in our thinking, we often do not explicitly consider or clearly articulate the value tradeoffs. Our inadequate thinking about value tradeoffs is due to three major errors:

1. an incomplete understanding of all of our fundamental objectives.
2. not recognizing the value tradeoffs that exist.
3. an unwillingness to think hard about and make necessary value tradeoffs.

Incomplete Understanding of All the Fundamental Objectives. On November 1, 2001, the then Governor Gray Davis of California publicly announced that he had received information from federal authorities that terrorists may attempt to destroy one or more of the major bridges in the state. Governor Davis was criticized by many for his remarks based on information that others said was not credible or corroborated. In defending his public announcement, Governor Davis said, "If I didn't make that statement and, God forbid, something happened, I'd be kicking myself." Was Governor Davis' announcement the result of a good decision? That depends on many things including the specific information he had, the objectives he considered in deciding, and the value tradeoffs he made.

What are the fundamental objectives that Governor Davis should use for appraising alternatives about whether to make such a public announcement and exactly what announcement to make. From a public perspective, I think that the major objectives should include those in Table 1 concerning loss of life, injuries, property damage, disruption, fear and anxiety, and costs of responding. The announcement could have lowered the chance that the bridges were attacked then or ever, or it could have lowered the chance of attack then and increased chances in the future, and/or it could have increased or decreased the chances that other targets are pursued. Significant uncertainties about the consequences of each alternative render the choice very difficult.

What were the fundamental objectives that Governor Davis thought about while contemplating a public announcement? His defense of his announcement was to avoid possibly "kicking himself". But where on the list of important fundamental objectives is the desire to minimize Governor Davis' personal and political regret? Its relevance is insignificant compared to the importance of potential loss of life and the other possible fundamental consequences. How much did Governor Davis weigh his potential regret compared to these real consequences? One cannot judge from press reports, but I'd have more confidence in the quality of Governor Davis' decision if he had said something like, "I believe the consequences to the public, in terms of potential loss of life and disruption of society, are less given the announcement." It is essential to keep in mind a clear understanding of our fundamental objectives to even begin to make appropriate decisions.

Not Recognizing Value Tradeoffs that Exist. If we do not recognize the value tradeoffs that exist, it is not possible to make informed choices based on a clear understanding of the consequences of different alternatives. Decisions involving terrorism evoke personal value tradeoffs among competing objectives, value tradeoffs between individuals, individual versus societal tradeoffs, and tradeoffs between different societal objectives.

Aware of a threat of bridge destruction, an individual must balance his or her loss of freedom and inconvenience of not using the bridge against the potential safety consequences of using it. The state of California on November 11, 2001 stopped and

searched all large trucks before allowing them to cross the Golden Gate Bridge. This inconvenienced and reduced the freedom of truckers and potentially increased the safety of all bridge users.

When we speak of civil liberties, it is important to clarify whose civil liberties are of concern. Subsequent to September 2001, the State Department began investigating backgrounds more thoroughly before granting visas to young men from Arab and Muslim nations. Pro-immigration groups criticized this decision as one that reduces civil liberties and is antithetical to American values. To better understand the issue, it is useful to distinguish between civil liberties of Americans and civil liberties of others seeking to come to the US. One cherished American civil liberty is the freedom to go where we want and do what we want when it does not harm others. Terrorism and potential terrorism limit this freedom. A more thorough control of visas may prevent some terrorism, and thereby enhance civil liberty for Americans and permanent residents, while it reduces a civil liberty to some visa applicants. This is a value tradeoff that cannot be avoided. To decide which competing value should be stressed here depends on the relative values for who and how many people are affected in what ways by the alternatives—in this case, reducing fear and increasing freedom for millions of Americans versus requiring a more thorough visa investigation for selected applicants.

Unwillingness to Make Value Tradeoffs. We have all heard people claim that a specific right is so important that we cannot or should not give up any amount, however small, of that right regardless of the other consequences. Some people claim that freedom of speech or freedom of the press are paramount to other values we all cherish. But these values do conflict with other critical values, such as life, liberty, and the pursuit of happiness. Choices must be made that involve tradeoffs among these cherished values.

You cannot yell "fire" in a crowded theatre when no fire exists. This could lead to deaths and injuries in an ensuing panic. The value tradeoff made is that one's freedom of speech to yell fire is not as significant as the freedom to live of many others. Such value tradeoffs clearly are recognized and made. Is this not the same type of value issue in examining the freedom of the press to announce planned American military action,

details about steps to apprehend terrorists, or specific vulnerabilities of infrastructure to terrorism? Balancing freedom of the press to release such information against the safety consequences to the public is very difficult. It must be made on a case by case basis. But one important point is that such value tradeoffs must be made in these situations. Given this, a second important point is that these value tradeoffs should be consistent with each other and made based on an understanding of the consequences and hard thinking about the values.

The Justice Department recently decided to allow authorities to monitor all communication between some people in federal custody and their lawyers. The stated purpose is to prevent possible future terrorist acts to save American lives. The President of the American Bar Association reportedly said, "No privilege is more indelibly ensconced in the American legal system than the attorney-client privilege." That may be accurate, but explicit value tradeoffs should still be made to consider the relative merits of this Justice Department action.

One of the detained individuals whose communication is to be monitored is Omar Abdul Rahman, who was convicted of the 1993 World Trade Center bombing. What are the appropriate value tradeoffs between the number of American lives that may be lost in the future due to terrorism, the number of individuals in custody whose communications are monitored, and any precedence set by the choices? Failure to address these value tradeoffs explicitly cannot lead to an informed choice.

Thinking About Value Tradeoffs. An appropriate value model for a complex decision specifies all the value tradeoffs necessary to balance the pros and cons of the alternatives. Consider the case described above concerning certain individuals in federal custody and their lawyers. The fundamental objectives include a desire to reduce loss of life due to terrorism, reduce lack of freedom of movement, and maintain the civil liberty of private attorney-client discussions.

To evaluate the pros and cons of whether such a ruling is desirable, we should estimate its consequences as well as we can. After a careful examination, suppose that the

monitoring of communications was felt to result in a ten percent chance of avoiding a terrorist attack that would kill 50 people, about twenty million people feeling more comfortable to go where they please safely, one million people who would not previously fly would be willing to do so, 1200 people in federal custody would have their conversations with lawyers monitored, and a possible precedence to limit the attorney-client privilege.

Suppose the value model is the additive utility function (1). After assessing an appropriate function it may indicate that the loss of attorney-client privilege for 500 people in this situation is equivalently as bad as the loss of one individual's life to a terrorist attack. Then the positive consequence of avoiding the fatalities due to terrorism outweighs the negative consequence of 1200 people losing their attorney-client privilege. The precedence set by monitoring conversations in this case may be viewed as being as significant (meaning bad) as having five million people feeling uncomfortable about moving around. Given this value tradeoff, the positive consequence of 20 million people feeling more free would outweigh the negative consequence of setting precedence for suspending attorney-client privilege in certain situations.

Appraisal of DHS Policies. With a decision analysis model such as illustrated in Figure 1 and a DHS value model as discussed, we could appraise different policies being considered. For example, we could examine the appropriateness of spending several million dollars to harden a particular stationary target, such as a nuclear power plant or a major dam. We could presume that this would reduce the likelihood of a successful terrorist attempt on the hardened target. If the terrorists knew that a particular target was hardened, they could divert their attention to different nuclear power plants or major dams. The probabilities of those alternate targets might be slightly increased. Hence, it may be concluded that investing in hardening stationary targets, if there were similar available targets, would have very little positive effect from the national viewpoint and yet utilize significant national resources. This type of understanding might be developed by using a value model in analysis to represent an intelligent adversary that can take actions to thwart DHS actions. Such an analysis is closely related to the interdependent security problem examined by Kunreuther and Heal (2003).

Reducing Terrorist Success by Reducing Consequences. Analysis may indicate that for some classes of terrorist acts, the indirect consequences of those acts are more detrimental to our nation than the direct consequences. This may be the case even though the direct consequences could be substantial. An analysis of consequences of such terrorist acts should lend some insight into the mechanisms by which the indirect consequences occur. For instance, as a result of the terrorist attacks of September 2001, our national airline system was shut down for about a week and the level of airline travel was significantly lower over the next few years. The indirect consequences this had on the economy, jobs, and businesses was significant. The impact in terms of fear and unhappiness was also significantly great. It may be that better and more accurate communication delivered to the public in a more timely manner would alleviate some of those indirect consequences. The value model in the analysis facilitates the creation of alternatives in these situations.

Communicating About Decisions. When decisions are made, they often need to be presented to others for justification and to garner support for the choice. As the choice is made to best achieve the objectives, discussing how and why that choice achieves those objectives well is a way to gain support for the decision. It also provides a basis to motivate others who need to take actions to implement the decision. If they believe the choice was made well and was founded on sound logic and the interests of our nation, it is likely that individuals will do their best to implement it. If they feel that the choice was poorly made, will be ineffectual, and/or is based on political considerations, the likely follow-through of managers and the public to ensure that its intended consequences occur, will likely not be fulfilled. A quality DHS value model founded on an agreed upon set of fundamental objectives provides a solid basis for this communication.

Prioritizing Domestic Targets from a Terrorist Perspective. If we have a utility function for possible terrorist activities, we can use this to help prioritize domestic targets from the terrorists' viewpoints. For this, we would need to estimate the likelihood that any particular attack on a target would be successful if attempted. Conditional on its success, we would want to estimate the possible consequences in terms of the terrorist

objectives. We'd also want to know what the consequences of failure would be in terms of their objectives. Then, using the value model, we could calculate the expected utility of any particular target. Those targets with the highest expected utility would be preferred if the prescriptive model were appropriate. Obviously for such a use, we would not know precisely the terrorist utility function. Hence, several sensitivity analyses would be appropriate that varied the terrorist value tradeoffs and perhaps their risk attitudes. We would likely find out that some targets seemed much more desirable than others and why for a wide range of potential terrorist values. This information could be useful in our planning.

Evaluating Terrorist Actions. For any particular terrorist target, an analysis of whether or not a terrorist action should be undertaken, and if so how it might be attempted, would be insightful. Of course, the caveat of using a prescriptive value model to understand possible terrorist actions needs to be taken into account when interpreting such an analysis. Still, when used with sensitivity analysis, some important insights could be developed. The target of plutonium theft discussed in section 4 presents an example for discussing the evaluating possible terrorist actions.

It is likely not reasonable to assume a terrorist would necessarily attempt to steal plutonium with the highest utility. More likely, this information would be combined with judgments of the likely success of achieving each of the possible targets in selecting a target. The proposed model suggests the terrorists would select the target with the highest expected utility.

A standard procedure to describe possible terrorist actions involves the generation of scenarios (i.e. terrorist acts). To obtain plutonium, a team of analysts can generate a large number of scenarios, some of which may turn out to be very desirable to the terrorist and yet have otherwise been previously overlooked in DHS planning. Using the terrorist value model and judgments about the likely "success" of various actions, an analysis can evaluate the relative desirability of the acts from the terrorist perspective. Because of government's uncertainties about the adversaries' preferences and judgments, sensitivity analyses incorporating an expected utility adversary model could be used to produce

probabilities of various adversary actions. These, along with DHS's value model, would be used to evaluate alternatives for improving control systems in a model such as that in Figure 1.

6. Summary and Recommendations.

Values are critical to decision making, and our choice of alternatives is the only way we have to influence terrorism and its consequences. A value model that clearly articulates and prioritizes our objectives is the guide for each of those decisions. It is a worthwhile endeavor to construct such value models. If we don't know what it is we are trying to achieve, it doesn't make sense to spend billions of dollars trying to achieve it. It does make sense to spend hundreds of thousands of dollars trying to clearly articulate what it is that we hope to achieve. Then the subsequent billions can perhaps much better achieve it.

Building DHS value models relies on our relative strengths with respect to terrorism. Our decisions about which alternatives to choose depend on the likely consequences of those alternatives and our relative desirability or undesirability for those consequences. We won't have and can't have basic data on possible terrorist acts and their possible consequences. We cannot have a precise understanding of terrorist actions and what their consequences will be. We do have access to all the information necessary to construct DHS value models. All the information is in the minds of people at DHS and elsewhere knowledgeable about our national interests and concerns about terrorism.

Value models are also an effective way to spend knowledge resources of our country. We need one value model to evaluate numerous alternatives. For each of those alternatives, we need a separate analysis to estimate what the possible consequences will be. Thus, for any complex decision laid out as in Figure 1 where there are only two stages of DHS decisions, there could literally be thousands of alternatives, yet there would only be one value model to guide both the creation and evaluation of those alternatives.

We should have a dedicated effort by DHS or sponsored by DHS to develop quality value models for our counter-terrorism activities. This would likely pay for itself economically and in many other ways several times over. We should also have activities that develop value models to represent terrorist objectives. Value models provide guidance for the highest level of DHS and our nation. The tasks of developing those value models should be clearly assigned within DHS to a project with no other priorities. It is that important to have the tasks done well.

References

Kahneman, D. and A. Tversky (1979), Prospect Theory: An Analysis of Decision Under Risk, Econometrica, 47, 263-91.

Kahneman, D., P. Slovic, and A. Tversky, eds. (1982), Judgment Under Uncertainty: Heuristics and Biases, Cambridge University Press, Cambridge, England.

Keeney, R.L. (1977), A Preliminary Model of Adversary Preferences, Proceedings of the Lawrence Symposium on System and Decision Sciences, Western Periodicals, North Hollywood, California, 187-193.

Keeney, R.L. (1988), Building Models of Values, European Journal of Operational Research, 37, 149-157.

Keeney, R.L. (1992), Value Focused Thinking: A Path to Creative Decision Making, Harvard University Press, Cambridge, Massachusetts.

Keeney, R.L. (2001), Countering Terrorism: The Clash of Values, ORMS Today, 28, No. 6, 20-22.

Keeney, R.L. (2006), Developing Objectives and Attributes, Advances in Decision Analysis, W. Edwards, R.F. Miles, Jr., and D. von Winterfeldt, eds., Cambridge University Press, New York.

Keeney, R.L. and R.S. Gregory (2005), Selecting Attributes to Measure the Achievement of Objectives, Operations Research, 53, 1-11.

Keeney, R.L., and H. Raiffa (1976), Decisions with Multiple Objectives, Wiley, New York. Reprinted by Cambridge University Press, 1993.

Krantz, D.H., R.D. Luce, P. Suppes, and A. Tversky (1971), Foundations of Measurement, Vol. 1. Academic Press, New York.

Kunreuther, H. and G. Heal (2003), Interdependent Security, Journal of Risk and Uncertainty, 26, 231-249.

Luce, R.D., and H. Raiffa (1957), Games and Decisions, Wiley, New York.

Pratt, J.W. (1964), Risk Aversion in the Small and in the Large, Econometrica, 32, 22-34.

Raiffa, H. (1982), The Art and Science of Negotiation, Harvard University Press, Cambridge, Massachusetts.

Viscusi, W.K. and R.J. Zeckhauser (2003), Sacrificing Civil Liberties to Reduce Terrorism Risks, Journal of Risk and Uncertainty, 26, 99-120.

von Neumann, J., and O. Morgenstern (1947), Theory of Games and Economic Behavior, 2nd edition, Princeton University Press, Princeton, New Jersey.

von Winterfeldt, D. and W. Edwards (1986), Decision Analysis and Behavioral Research, Cambridge University Press, Cambridge, England.

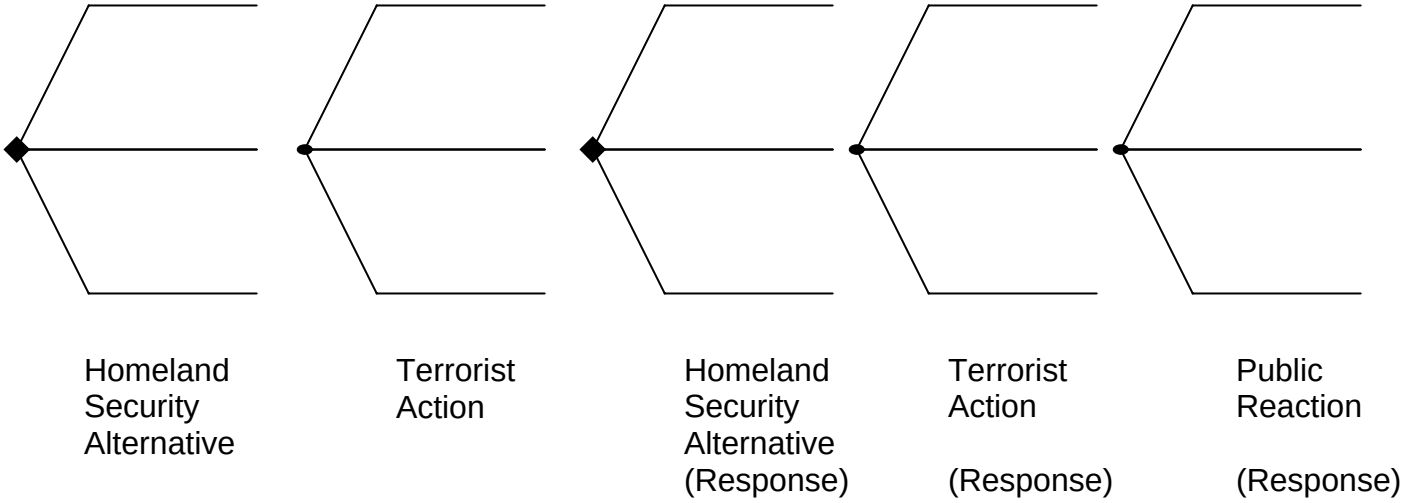

Figure 1. Decisions of Homeland Security with Terrorist and Public Action

Table 1. Preliminary Objectives and Attributes for Homeland Security

Objectives	**Attribute**
1. Minimize personal health and safety losses	
Minimize loss of life	Number of deaths
Minimize injuries	Number of disabled
2. Minimize economic damage	
Minimize personal economic damage	Cost in dollars
Minimize business economic damage	Cost in dollars
Minimize government economic damage	Cost in dollars
3. Minimize future terrorism	Constructed scale with positive, neutral, negative effects on recruiting terrorists
4. Minimize support for terrorism	
Minimize political support	National or international bodies publicly against US actions
Minimize economic support	Dollars flowing to terrorist causes
5. Minimize personal loss of freedom	Number of citizens experiencing important limits to their way of life (not flying, not shopping in malls)
6. Minimize fear and despair	Number of citizens experiencing significant fear or despair
7. Minimize indirect economic consequences of terrorism	
Minimize indirect economic consequences to individuals	Number of jobs lost
Minimize indirect consequences to business	Cost in dollars
Minimize indirect consequences to government	Cost in dollars
8. Minimize counter terrorism cost	Cost in dollars

Table 2. Ranges for the Terrorist Attributes

Attribute	Measure	Worst	Best
X_1 = Plutonium	g.Pu.	10	2500
X_2 = Purity	g. U./L.	333	0
X_3 = Radiation	Rads./hr.	100	0

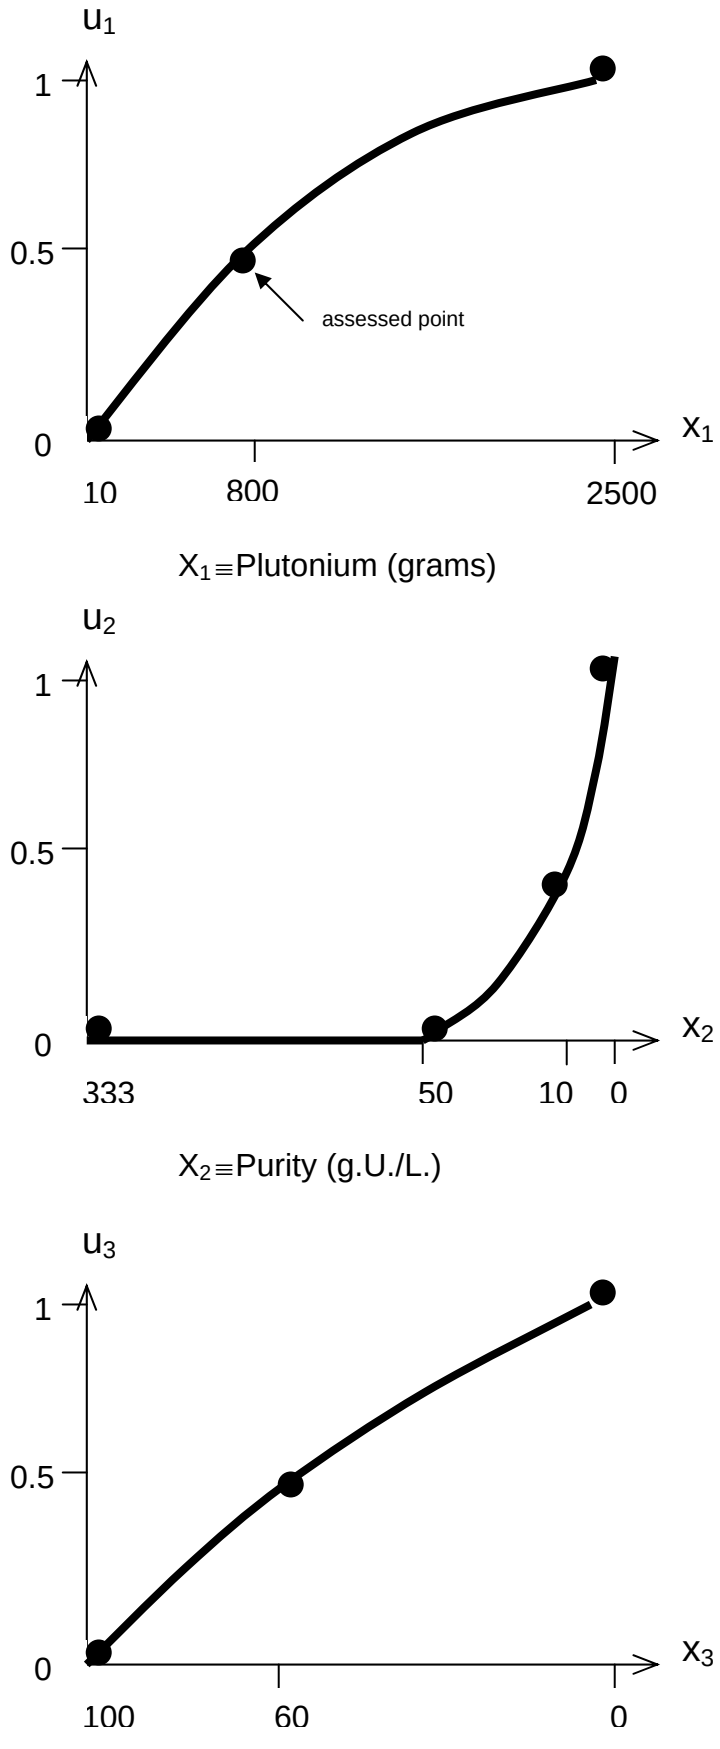

Figure 2. The Single-Attribute Utility Functions for Terrorist Value Model

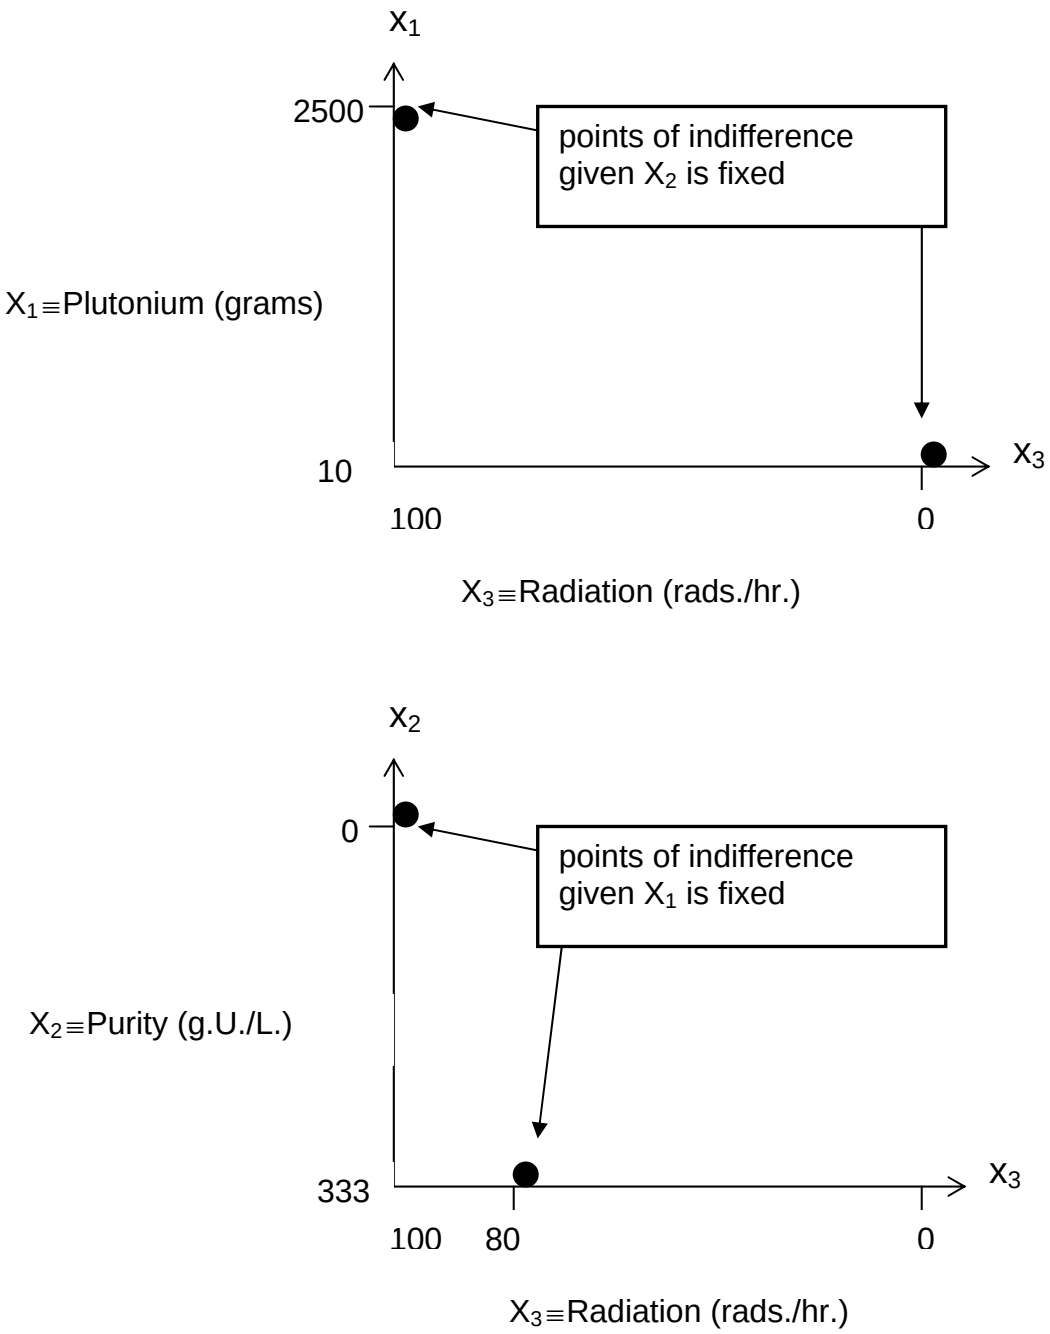

Figure 3. Value Tradeoffs for Terrorist Value Model