ANL/DIS-08/2

Influence of Time-dependent Factors in the Evaluation of Critical Infrastructure Protection Measures

Decision and Information Sciences Division

ANL/DIS-08/2

Influence of Time-dependent Factors in the Evaluation of Critical Infrastructure Protection Measures

by
W.A. Buehring and M.E. Samsa
Decision and Information Sciences Division, Argonne National Laboratory

for
U.S. Department of Homeland Security
Science and Technology Directorate

March 2008

CONTENTS

CONTENTS (Cont.)

FIGURES

FIGURES (Cont.)

TABLES

INFLUENCE OF TIME-DEPENDENT FACTORS IN THE EVALUATION OF CRITICAL INFRASTRUCTURE PROTECTION MEASURES

by

W.A. Buehring and M.E. Samsa

ABSTRACT

The examination of which protective measures are the most appropriate to be implemented in order to prevent, protect against, respond to, and recover from attacks on critical infrastructures and key resources typically involves a comparison of the consequences that could occur when the protective measure is implemented to those that could occur when it is not. This report describes a framework for evaluation that provides some additional capabilities for comparing optional protective measures. It illustrates some potentially important time-dependent factors, such as the implementation rate, that affect the relative pros and cons associated with widespread implementation of protective measures. It presents example results from the use of protective measures, such as detectors and pretrained responders, for an illustrative biological incident. Results show that the choice of an alternative measure can depend on whether or not policy and financial support can be maintained for extended periods of time. Choice of a time horizon greatly influences the comparison of alternatives.

SUMMARY

An examination of which protective measures would be the most appropriate to implement in order to prevent, protect against, respond to, and recover from attacks on critical infrastructures and key resources typically involves a comparison of the consequences of two case studies — one in which the protective measure was implemented, and one in which it was not. This report illustrates, through hypothetical examples, some potentially important time-dependent factors (e.g., implementation rate and recurring costs) that could affect the relative pros and cons associated with widespread implementation of protective measures.

The comparison of a set of scenarios in which the protective measure of interest is completely implemented to a set of scenarios in which it is not is customarily accomplished by considering the present value of implementing the protective measure and its operational costs. The question examined here is whether the consideration of a more refined representation of just a few key time-dependent factors could make a significant difference in the evaluation of the overall merit of the measure.

Time-dependent factors have been shown to affect the competitiveness of alternative protective measures. A framework that provides some additional factors to consider when evaluating alternative measures is described here. Although the framework draws extensively on the estimates of consequences taken from the Critical Infrastructure Protection Decision Support System (CIPDSS) model used by the U.S. Department of Homeland Security (DHS), it can be applied to any evaluation that employs estimates of consequences from implementing alternative protective measures.

The factors include:

1. Estimated threat levels that vary with time;

2. Time it takes to implement protective measures;

3. Partial capability of measures during implementation;

4. Causes of delays in making implementation decisions (e.g., funding limits);

5. Causes of delays in implementation (e.g., physical problems);

6. Lifetimes of equipment, training, and medicines;

7. Likelihood of breakeven incidents and time it takes to break even;

8. Yearly investment and operational costs and their compatibility with the DHS budget process (vs. total present value); and

9. Readiness of the framework for discounting if desired.

The hypothetical examples given in this report are intended to illustrate the importance of the first eight factors listed above and, in order to avoid further complexity, do not include the final factor of readiness for discounting.

This report presents the results from using various protective measures, including detectors and pretrained responders, to address an illustrative biological incident. The results provide insight into how long policies must be supported in order to result in overall benefits. Some of the examples show that a different protective measure alternative should be selected if policy and financial support for the measure in question cannot be maintained for more than 10 years. Some examples also show that although substantial overall benefits are possible, they may not appear until a number of years after the measure's initial implementation. The magnitude of the results and the relative desirability of the outcomes could change significantly with different assumptions about costs and threats.

1 INTRODUCTION

The examination of which optional measures[1] are the most appropriate to implement in order to prevent, protect against, respond to, and recover from attacks on critical infrastructures and key resources (CIKR) typically involves a comparison of the consequences from two case studies — one in which the protective measure is implemented and one in which it is not. This approach is generally used in studies that apply the Critical Infrastructure Protection Decision Support System (CIPDSS) model for the U.S. Department of Homeland Security (DHS). This report illustrates some potentially important time-dependent factors (e.g., implementation rate and recurring costs) that could affect the relative pros and cons associated with widespread implementation of protective measures.

1 Optional measures include a range of preparedness activities (e.g., planning, organizing, obtaining equipment, training, conducting exercises, managing, developing and choosing policy options) (DHS 2007). For simplicity, in this report, the combination of all these activities is referred to as "protective measures."

2 RISK AND THE CIPDSS FRAMEWORK

An evaluation of the pros and cons of implementing a protective measure depends on the framework in which the risk is estimated. This section briefly outlines the overall concept of risk and introduces how time-dependent factors could affect the deliberation over which measures to employ.

2.1 OVERALL RISK

The DHS risk management framework presented in the National Infrastructure Protection Plan (NIPP) is shown in Figure 2-1 (DHS 2006). "Risk" is generally defined as the combination of the frequency of occurrence, vulnerability, and consequence of a specified hazardous event. In the context of NIPP, risk is the expected magnitude of loss (e.g., deaths, injuries, economic damage, loss of public confidence, and loss of government capability) due to a terrorist attack, natural disaster, or other incident, along with the likelihood of such an event occurring and causing that loss. The NIPP risk management framework establishes the process for combining consequence, vulnerability, and threat information to produce a comprehensive, systematic, and rational assessment of national or sector-specific risk that drives CIKR protection activities. The framework applies to the general threat environment, as well as to specific threats or incident situations.

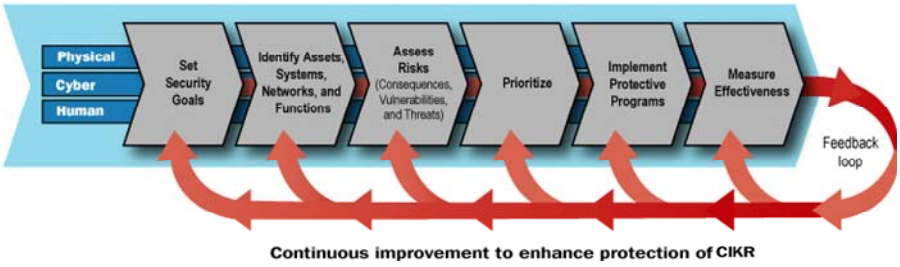

FIGURE 2-1 NIPP Risk Management Framework

Thus, protective measures that affect threats, vulnerabilities, and/or consequences have the potential to reduce risk. The key question is whether the risk reduction associated with implementing the measure is worth the cost. Estimating the effects of specific protective measures (including the effects of dependencies, interdependencies, and cascading events) to improve risk-informed decision making is a reason the simulation model described in the following section has been developed.

2.2 RELATIONSHIP BETWEEN RISK AND EXPECTED UTILITY

The term "risk" as commonly used within DHS is defined as

$$\text{Risk} = \text{Threat} \times \text{Vulnerability} \times \text{Consequences}.$$

Because "threat" is defined as the probability that an attack will be attempted, and "vulnerability" is defined as the probability that the attack will be successful if it is attempted, the risk equation can be restated as follows:

$$\text{Risk} = (\text{Probability of Attack}) \times (\text{Probability of Success, Given an Attack}) \times (\text{Consequences, Given an Attack}).$$

This equation and the following discussion can be simplified as follows: If "incident" is defined as a successful attack, then

$$\text{Risk} = (\text{Probability of an Incident}) \times (\text{Consequences, Given an Incident}).$$

This relationship seems simple enough, but for any given incident, the range of consequences can be great and varied. Consequences can include human fatalities and injuries; response, repair, and restoration costs; lost economic productivity (e.g., lost gross domestic product [GDP]); impacts on public confidence and national security; and other undesirable impacts on things that we as Americans value most. Discussions about the risk of terrorist attacks or natural disasters seldom go beyond the general terms in the above equation; thus, they rarely focus on the relative importance of each type of consequence or the decision maker's propensity to tolerate or avoid uncertainty, which is the basis for risk.

The construct of "expected utility" enhances and extends this concept of risk by combining the different types of consequences into a single "equivalent consequence" that is based on the decision maker's value structure. The construct then makes further adjustments on the basis of the decision maker's attitude toward uncertainty so that this single metric can be used to compare alternative options regardless of the degree (i.e., amount) of uncertainty in each option. The complete formulation includes the costs of any protective measures being implemented and the likelihood of there being no incident.

First, through proper elicitation of the decision maker, a value structure similar to that shown in Table 2-1 can be constructed. In this example, the decision maker's equivalent value for each consequence is specified in an economic metric (i.e., millions of dollars). Showing the equivalencies in any consequence unit (e.g., equivalent fatalities, equivalent injuries, or equivalent population confidence points) is straightforward.

Monte Carlo simulation models are typically used to calculate the consequences of a hypothesized terrorist attack or natural disaster because they are useful for exploring uncertainty in the resultant consequences. Each of the consequences for each of the simulation runs is multiplied by its respective equivalent value (which is based on the decision maker's value

TABLE 2-1 Example of a Decision Maker's Value Structure

Consequence	Consequence Amount (natural units)	Decision Maker's Equivalent Value (economic measure)
Cost of mitigation measures	$1 million	$1 million
Lost GDP	$1 million	$1 million
Response, repair, and restoration costs	$1 million	$1 million
Statistical fatalities	1 fatality	$10 million
Nonfatal moderate to serious injuries	1 injury	$0.2 million
Moderate to serious illnesses	1 illness	$0.2 million
Lost public confidence	1 million population confidence points	$50 million

structure) in order to generate (in this example) a combined equivalent cost for each simulation run. The combined equivalent cost is a single measure of consequence for each simulation run that explicitly accounts for the decision maker's value structure. Figure 2-2 is an example of a combined equivalent cost probability distribution produced in this manner. This type of chart is often considered a representation of risk, because it shows the probability distribution of an aggregate representation of undesirable outcomes. However, if a mitigation measure is under consideration, event consequences alone do not account for that cost, nor for the probability that no incident will occur. There is often uncertainty in the cost of the mitigation measure as well, and an example of this is represented in the protective measure cost distribution in Figure 2-3.

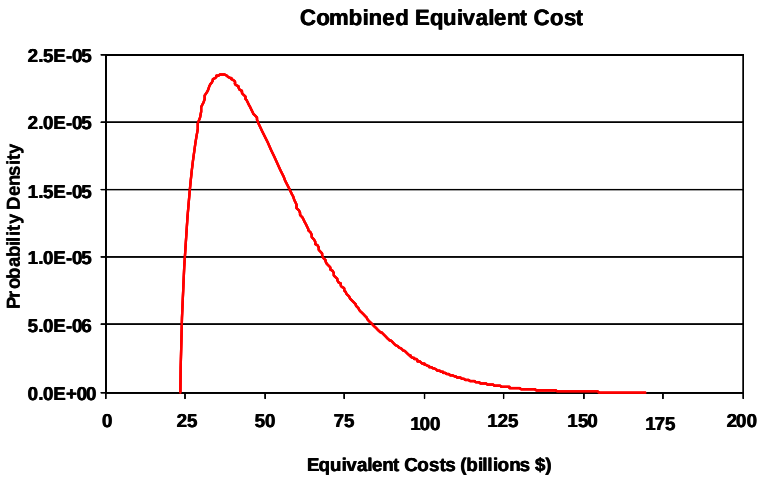

FIGURE 2-2 Example of Equivalent Cost Distribution

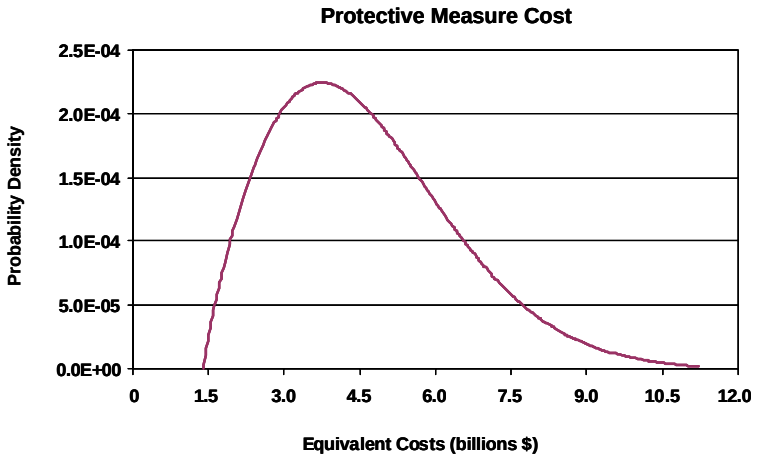

FIGURE 2-3 Example of Protective Measure Cost Distribution

Because individual decision makers have different attitudes toward uncertainty, they interpret probability distributions like these in different ways and draw different conclusions from them. A decision maker who is characterized as being "neutral" toward uncertainty (and, therefore, risk) would tend to choose between alternatives on the basis of the average, or expected, value of the uncertain outcomes. A risk-averse decision maker would rather choose a known or certain outcome that is greater than the expected outcome of the uncertainty distribution in the hope of avoiding an uncertain outcome that is even worse than the cost he/she would elect to accept to avoid future uncertainty. For example, many of us buy insurance because we are averse to risk. By contrast, a risk-tolerant decision maker would accept a certain cost that is less than the expected value of the uncertain outcomes in the hope of getting an outcome that is better the expected outcome. These characteristics define the decision maker's "utility" function, or the relative preference he/she has for each level of possible outcome. Three prototypical utility functions are shown in Figure 2-4.

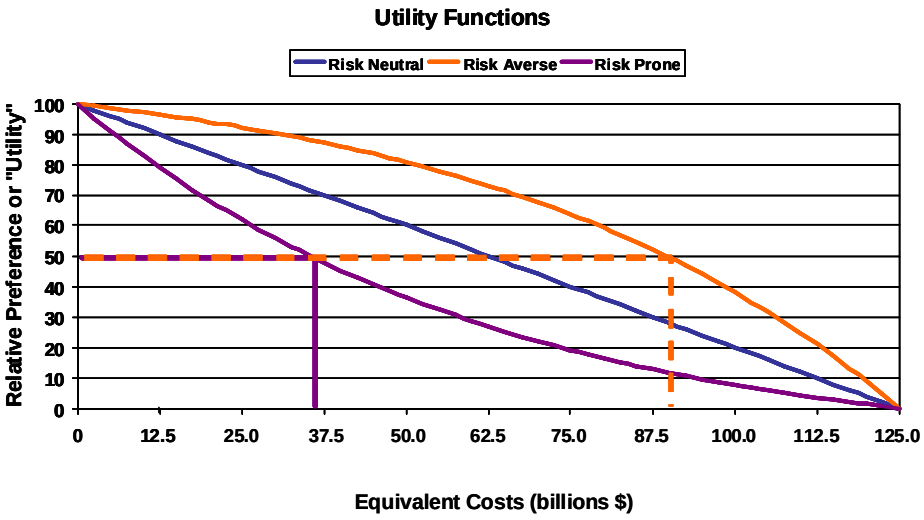

FIGURE 2-4 Prototypical Decision Maker Utility for Equivalent Cost Functions

By eliciting the decision maker's utility for equivalent costs and using this function to convert each simulated consequence outcome to a relative preference (i.e., utility value), one accounts for any nonlinearity in the decision maker's attitude toward uncertainty and risk and converts the simulated consequences to an equivalent linear utility scale. The average of this linear utility measure (i.e., expected utility) is thus a single meaningful metric of relative risk that fully incorporates the decision maker's value structure for different consequences and his attitude toward uncertain outcomes.

Risk — or, more precisely, relative risk, because the utility scale here has been arbitrarily set between 0 and 100 — is thus given by the following:

$$\text{Relative Risk} = P \times (\text{Expected Utility for Equivalent Costs}) + (1 - P) \times (\text{Expected Utility for the Protective Measure Cost})$$

where P = likelihood of an incident and (1 – P) = likelihood of no incident.

For decision makers who are more comfortable thinking in terms of consequence levels rather than utility, it is a simple matter to convert the utility scale to a equivalent consequence measure by multiplying the utility value by the appropriate conversion vector.

2.3 CRITICAL INFRASTRUCTURE PROTECTION DECISION SUPPORT SYSTEM

The CIPDSS is a support tool for making risk-informed decisions on the basis of system dynamics models that estimate the interdependent consequences in 17 (an 18th is being considered) U.S. infrastructures that would result from a disruptive event, such as a terrorist attack. These consequence models are coupled to a multiattribute decision model that combines uncertain outcomes with the value trade-offs and risk attitudes of the decision-making organization into a single relative preference measure (i.e., utility) for each alternative or decision option. This single relative preference measure facilitates the comparison of alternative protective measures in situations where there may be a lot of uncertainty about the effectiveness of the alternatives and where the choices are based on multiple and often conflicting objectives. The CIPDSS decision metrics include economic, human health and safety, environmental, sociopolitical, and national security considerations.

A value structure (i.e., set of value trade-offs), which measures the relative importance of each decision metric, and a risk attitude function, which translates outcomes drawn for uncertain distributions into the decision-making organization's "certainty equivalents," make up a decision-maker profile. The CIPDSS methodology reflects the characteristics and ranges of value trade-offs and risk attitudes elicited from a number of DHS decision makers and other experts.

Within the CIPDSS framework, the decision metrics are generated by an extensive collection of interdependent infrastructure systems' dynamic models that simulate, over a specified time horizon, detailed impacts resulting from a hypothetical event or scenario. The CIPDSS simulations of incident scenarios produce an array of consequences that start at a time

just before the postulated incident and continue until the time that the effect of the incident has disappeared or a new equilibrium has been reached (which may be more than a year). The detailed impacts are then categorized and aggregated in the high-level decision metrics. A case-by-case approach is used to calculate and store the decision metrics associated with the set of alternative protective measures or decision choices (Figure 2-5). A case is composed of at least two scenario pairs: (1) one base scenario pair and (2) one or more alternate scenario pairs. Each scenario pair includes a readiness scenario and an incident scenario.

The base readiness scenario depicts the state of the world just before and during the period of the postulated event. Existing protective measures are in place, and no event is postulated. This scenario represents the status quo. Then an event, such as a terrorist attack, is assumed to be perpetrated on the base readiness state, resulting in the base incident scenario. The base incident scenario measures the consequences of the postulated event in the absence of any of the additional protective measures to be considered in the decision.

Next, one of the alternative protective measures or decision options being considered is described and characterized within the base readiness state, resulting in an alternative readiness scenario. Often, the only difference between the base readiness scenario and an alternative readiness scenario is the cost and operational impacts associated with the protective measure being considered (the cost of the protective measure is generally included as a present value of the cost stream associated with implementation and operation). Finally, the postulated event is again perpetrated on the alternative readiness state, resulting in an alternative incident scenario with consequences that would be mitigated or otherwise affected by the alternative protective measure, relative to the base incident scenario. A pair of alternative scenarios is constructed in like manner for each alternative protective measure under consideration.

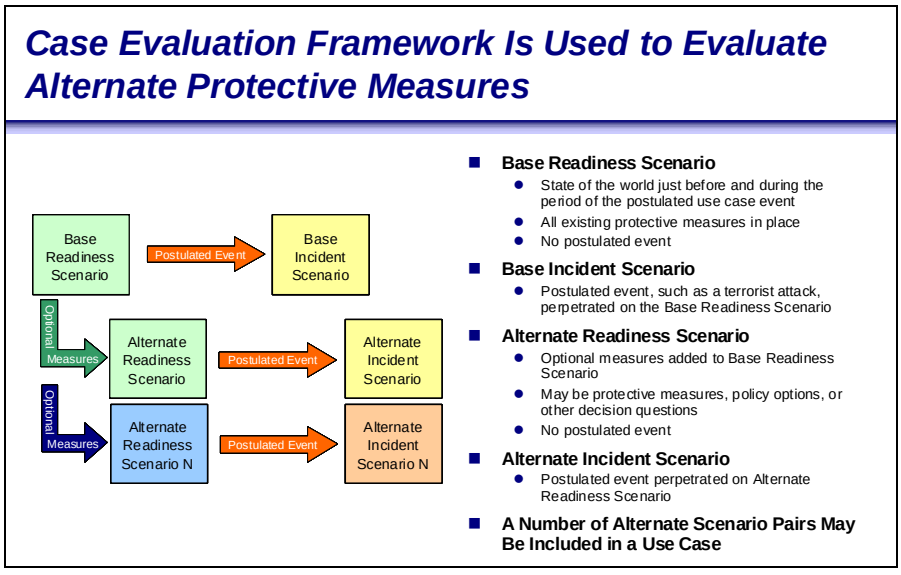

FIGURE 2-5 CIPDSS Evaluation Framework

The CIPDSS results or other consequence estimates can be used to estimate the minimum incident likelihood that justifies implementing a particular set of protective measures. For example, a base scenario and an alternate scenario that compares the no action case with implementation of a protective measure are analyzed as usual with CIPDSS (Figure 2-6). The reduction in consequences, as measured by CIPDSS decision metrics, between the alternate scenario and the base scenario is then compared with the investment and operation costs and any other impacts associated with protective measure implementation to determine the minimum incident likelihood at which this particular investment is justified. Any incident likelihood equal to or greater than this minimum would indicate that the alternate scenario investments are better than the base scenario conditions (no action).

The expected utility (i.e., relative preference) for each scenario can be plotted against the likelihood of an incident to determine which of several optional protective measures is best suited for implementation. The illustrative results in Figure 2-6 indicate that the preferred option depends on the likelihood of an incident. If intelligence (or judgment) suggests that the likelihood of an incident is the range labeled as low in Figure 2-6, the highest expected utility is associated with taking no additional action. If the likelihood of an incident is in the range labeled medium, the highest utility is associated with Alternate A. If the likelihood of an incident is above the medium range, investing in Alternate B has the highest utility.

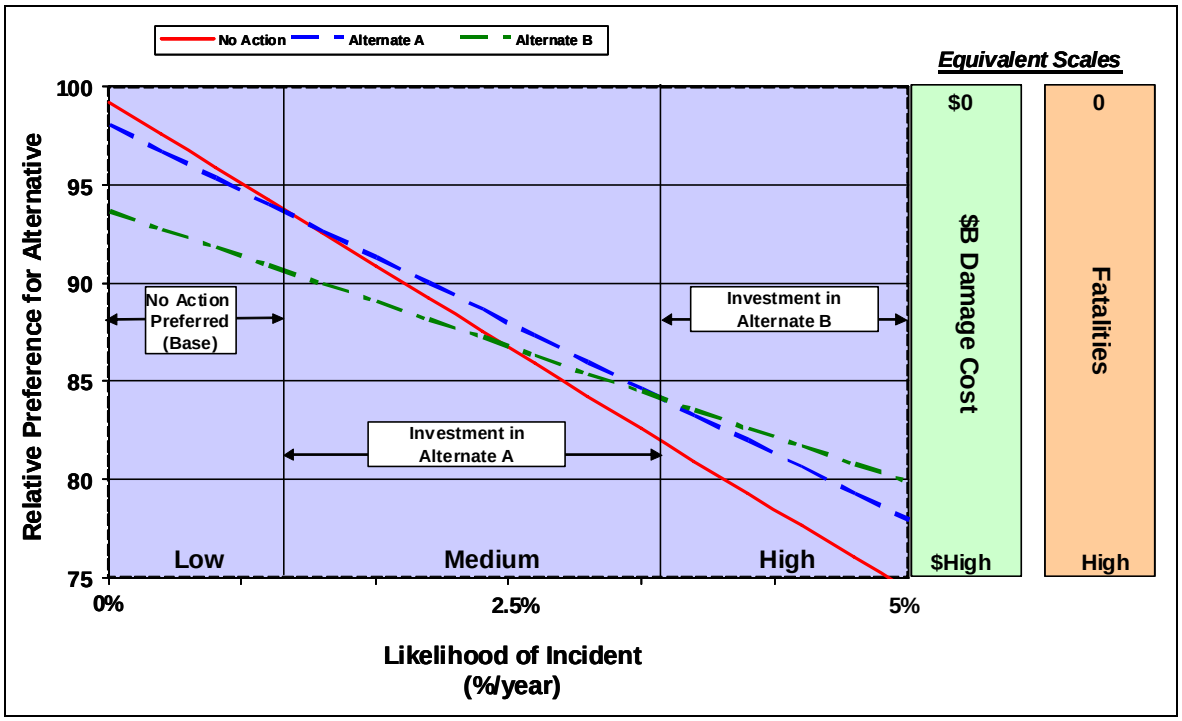

FIGURE 2-6 Illustrative CIPDSS Results

Often, relative preference or utility, which is typically normalized to a 0–1 or 0–100 scale, is not as meaningful to decision makers as a scale in one of the decision metrics. It is straightforward to translate utility back into any of the original decision metrics and show an equivalent scale, as illustrated by the right-hand scales for equivalent damage cost and equivalent fatalities. This conversion is sometimes helpful in understanding whether small differences in utility are important, as demonstrated later in this report.

The illustrative results shown in Figure 2-6 indicate that there are ranges of incident likelihood at which a particular option is clearly preferred, while there are other ranges at which the selection of the preferred option is a close call. For example, when the incident likelihood is zero, the best option is clearly the no action (no cost) option. As the likelihood of the incident increases, the no action option becomes less preferable than other options that have costs but also have benefits in terms of reduced risk, as measured by the decision metrics. Estimating the incident likelihood ranges at which a particular protective measure is clearly preferred and the incident likelihood ranges at which another protective measure or the no-action option is a strong competitor can be affected by the representation of time-dependent factors, such as those described in Section 3.

Since the investment in protective measures is made before the attack, the comparison ultimately comes down to risk reduction (measured by comparing CIPDSS outputs for the base scenario with the alternate scenario combined with the incident likelihood) versus expenditures and other impacts associated with protective measure implementation. By examining different investment strategies with CIPDSS (e.g., different protective measures or different levels of implementation or different locations for implementation), DHS can gather information on appropriate investment strategies for a given threat level.

3 TIME-DEPENDENT FACTORS FOR CONSIDERATION

The comparison of a set of scenarios (one in which the protective measure of interest is completely implemented and one in which it is not) is customarily accomplished by considering the present value of implementing the protective measure and its operational costs. The question examined here is whether the consideration of a more refined representation of a few key time-dependent factors could make a significant difference in the evaluation of the overall merit of the measure being considered. The key potential time-dependent factors associated with the protective measures are briefly described in this section.

3.1 MINIMUM TIME NEEDED TO ATTAIN THE MEASURE'S CAPABILITY

The broad implementation of a protective measure (e.g., using biological detectors or pretraining medical responders) would likely take a number of years, once the decision to proceed has been made. For example, in the illustrative case involving biodetectors presented in Section 4, it would take 4 years to fully install the detectors.

3.2 DELAYS IN DECIDING OR DURING IMPLEMENTATION

Once the decision to implement a particular protective measure has been made, a number of factors could affect the planned rate at which it is implemented. These include the need to make budgetary decisions, the use of new technology or information that affects the performance of the measure, and complications in procuring and implementing and the measure. In general, it can be assumed that these factors would cause delays in implementation; however, it is also possible that some factors could accelerate implementation, and these should also be analyzed.

3.3 LIFETIME AND MAINTENANCE OF THE PROTECTIVE MEASURE

Protective measures (e.g., detectors) have finite lifetimes and need replacement. Items of equipment (same example: detectors) also need maintenance periodically to stay in top working order. The training of responders takes a few years, and the training needs to be repeated from time to time as new responders join and as veterans need their training refreshed. Special medicines, such as anti-viral drugs, have a shelf lifetime and need replacement. It would be an unusual situation if the protective measure lifetime would correspond precisely with the time horizon used in the simulation of consequences.

3.4 POTENTIAL CHANGING THREAT LEVELS

The likelihood of an incident occurring is linked to the estimated threat levels, which may change over time for a number of reasons. New intelligence may become available. Terrorist organizations are likely to enhance their capabilities over time. Terrorists may pursue some types

of threats more than others because of the relatively "favorable" consequences that could result from such an attack. For example, a biological attack may be a preferred threat that is currently not likely to occur because of a lack of capability but that could become more likely because the potential consequences appear so desirable. Changes over time of the likelihood of *any* attack as well as changes in the relative likelihoods of the various types of threats must be recognized.

3.5 EFFECT OF THE PROTECTIVE MEASURE ON THE THREAT AND VULNERABILITY

Analyzing the effects of protective measures with models such as CIPDSS yields an estimate of the changes in consequences that would result from implementing the protective measure. Most protective measures reduce vulnerability, and they may also reduce the threat. For example, changing the rules about carrying liquids onto airplanes presumably has greatly reduced the likelihood of an incident associated with that threat.

3.6 EFFECT OF IMPLEMENTING MORE THAN ONE OPTION AT DIFFERENT TIMES

If a number of protective measures are being considered as candidates to combat a particular threat, the best strategy might be to implement more than one measure at different times rather than implementing all of them at once. This strategy could be necessary because of budgetary restrictions, for example. A series of CIPDSS case studies could be used to simulate the effect of an incident when different combinations of protective measures are implemented. At the beginning of the time horizon, each CIPDSS case study would estimate the impacts of an incident for over a year or more for a specified level of each protective measure. Then a year-by-year analysis could be conducted by weighting the CIPDSS results by the annual likelihood of the incident (also potentially a function of time).

3.7 DISCOUNTING OF COSTS AND OTHER CONSEQUENCES

Although none of the costs or other consequences described in this report are discounted, the topic is worthy of consideration in the evaluation of alternative protective measures. Discounting is the most commonly used and understood time-dependent factor. Hypothetical examples given in this report are intended to illustrate the importance of the six factors listed above and do not include the seventh factor of discounting in order to avoid further complexity. Real-world evaluations of alternative protective measures should address discounting explicitly.

During the preference assessment that provides the decision-maker profile that is part of the CIPDSS decision analysis framework, a few questions can be asked to determine whether, and to what degree, the value of a unit of consequence in the first year differs from the value of an identical unit of consequence occurring in a later year. By asking questions about relative

value for each consequence separately, an *implicit discount rate*[2] can be determined for each consequence (e.g., fatalities, costs, public confidence, national security, and environment). The discounted consequence stream over time can then be converted to an equivalent level of consequence at the point in time used as a basis for discounting (e.g., a typical basis point in time might be the time of the incident). Then the reference utility function being used in CIPDSS can be applied to the set of equivalent consequences.

This approach accommodates the possibility that different implicit discount rates may be assessed for different consequences. For example, an annual discount rate of 3% may be used for monetary values, while a statistical fatality may have the same assessed value at all times (i.e., implicit discount rate for statistical fatalities is zero). The implications of discounting or not discounting statistical fatalities are addressed in the decision analysis literature; for example, see Keeney (1995).

Traditional cost-benefit analysis typically uses constant monetary values per unit of consequence (benefit or cost) to convert consequences to monetary terms over the entire range for each consequence. The equivalent monetary values are then discounted to a common point in time (using a single discount rate), and the net present value of the alternative is computed and compared to other alternatives.

In 1992, the White House Office of Management and Budget (OMB) published a document providing guidance for heads of executive departments on how to conduct a benefit-cost analysis of federal programs (OMB 1992):

> The standard criterion for deciding whether a government program can be justified on economic principles is *net present value* — the discounted monetized value of expected net benefits (i.e., benefits minus costs). Net present value is computed by assigning monetary values to benefits and costs, discounting future benefits and costs *using an appropriate discount rate* [emphasis added], and subtracting the sum total of discounted costs from the sum total of discounted benefits. Discounting benefits and costs transforms gains and losses occurring in different time periods to a common unit of measurement. Programs with a positive net present value increase social resources and are generally preferred. Programs with a negative net present value should generally be avoided. (Sec. 5a, *Net Present Value and Related Outcome Measures*)

---

2 Discounting consequences to a common point in time is a time-honored approach for addressing consequences occurring at different times (especially those valued in monetary terms). The discount rate is the interest rate used in calculating the present value of expected yearly consequences. The discount factor is the factor that translates consequences in any given future year into present value terms. The discount factor is equal to $1/(1 + i)^t$, where i is the annual interest rate and t is the number of years from the basis date for present value calculations to the time when the consequence to be discounted occurs.

The OMB circular also had recommendations about discounting nonmonetized benefits and costs. There was no question that time-dependent values should be included (OMB 1992):

> In order to compute net present value, it is necessary to discount future benefits and costs. This discounting reflects the time value of money. Benefits and costs are worth more if they are experienced sooner. All future benefits and costs, including nonmonetized benefits and costs, should be discounted. (Sec. 8: *Discount Rate Policy*)

The OMB circular also addressed uncertainty and stated that, in situations with uncertain consequences, the expected values of the distributions of consequences can be obtained by weighting each consequence by its probability of occurrence, then summing across all potential consequences. In general, the expected value was considered the appropriate estimate to use (Sec. 9b). However, the OMB circular also allowed for the possibility that risky situations may need special considerations:

> In general, variations in the discount rate are not the appropriate method of adjusting net present value for the special risks of particular projects. In some cases, it may be possible to estimate *certainty-equivalents* ... to account for risk. (Sec. 9d: *Other Adjustments for Uncertainty*)

The CIPDSS decision analysis framework uses this method to combine an outcome distribution for an uncertain consequence with an assessed risk attitude to yield a certainty equivalent. For a risk-neutral person, the certainty equivalent is the expected value of the outcome distribution for the consequence. For situations in which the decision-maker profile exhibits a risk attitude with respect to any consequence that is not risk neutral (i.e., the risk attitude is risk averse or risk tolerant), the certainty equivalent of the consequence distribution is *not* equal to the expected value of the consequence distribution for the consequence, and CIPDSS has the capability to recognize and appropriately account for such situations, in addition to the situations that are totally risk neutral. Furthermore, a risk attitude other than risk neutral (determined by assessing preferences under conditions of uncertainty) with respect to a consequence implies that the risk-adjusted monetary value per unit of consequence is not constant across the entire range.

4 REPRESENTATION OF OPTIONS FOR ILLUSTRATIVE BIOLOGICAL THREAT SCENARIOS

A simplified analysis that draws on consequence estimates, such as those that could be obtained from CIPDSS, is used to demonstrate the potential effect of considering selected time-dependent factors. The protective measure options for an illustrative biological threat scenario are outlined along with a representation of associated time-dependent threat levels.

A 25-year time frame is used for the examples in Section 6 to allow up to 5 years to build up to the full implementation of the protective measure and a period of at least 20 years with that full implementation in force. The expected consequences and associated utility for a given year are determined by multiplying the appropriate readiness scenario by the likelihood that the incident does not occur and by multiplying the appropriate incident scenario by the likelihood that the incident does occur (as in Figure 2-5).

4.1 EXISTING MEASURES

The basis for comparison, as depicted in Figure 2-5, is a postulated biological incident with no additional protective measures beyond those currently in place. The impacts of the incident are examined over a 25-year time horizon.

4.2 ALTERNATIVE A: BIODETECTORS

Placing biodetectors at various key locations may allow an earlier identification of whether an incident has occurred and what agent was used. This increased speed could reduce fatalities and result in less public response and associated economic impacts.

The following key assumptions were used for the illustrative detector scenario:

- Detectors are implemented in key locations in the most populous U.S. cities, in subways, railway tunnels, large airports, and large domed stadiums.
- The first detectors are purchased in year 1 and implemented in year 2.
- By the end of year 5, all detectors are implemented.
- The detector lifetime is 8 years.
- One-time fixed costs include the detector capital cost and installation cost over 4 years.
- Operating costs include periodic testing and sample collection and analysis costs.

For a 25-year analysis, the consequences, other than detector cost for the first year, are associated entirely with the existing measures case (no detectors are in place). For the second year, the consequences are estimated as 25% of the consequences for the biodetector case and 75% of the consequences for the existing measures case (25% of the total detectors are operating). The representation proceeds as shown in Table 4-1. Since the detectors are assumed to have an 8-year lifetime, the detectors placed in operation at the beginning of year 2 must be replaced by the beginning of year 10. The total annual costs over years 1 through 25 for implementing this option range from $41 million to $224 million, as shown in Figure 4-1.

TABLE 4-1 Biodetector Annual Representation

	Weighting of Consequence Case		No. of New Detectors Implemented (% of total at beginning of year)
Year	Existing Measures (%)	Biodetector (%)	
1	100	0	0
2	75	25	25
3	50	50	25
4	25	75	25
5	0	100	25
6	0	100	0
7	0	100	0
8	0	100	0
9	0	100	0
10	0	100	25
11	0	100	25
…	…	…	…
25	0	100	25

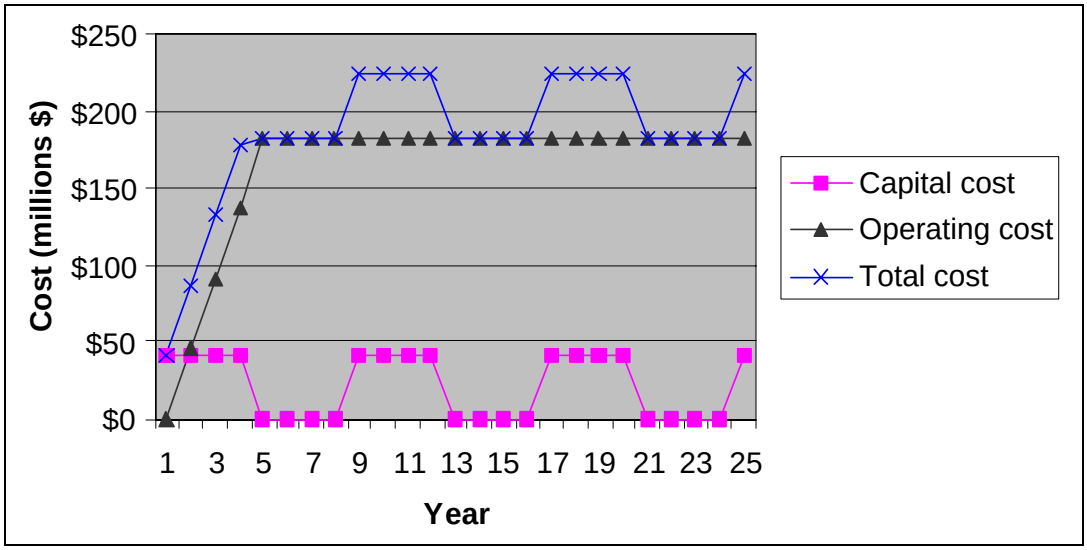

FIGURE 4-1 Annual Cost Distribution for the Biodetector Alternative

4.3 ALTERNATIVE B: ANTI-VIRAL DRUGS

The new anti-viral drugs developed and distributed throughout the United States could reduce the consequences from those assumed in the existing measures case. The following key assumptions were used for the illustrative anti-viral drugs scenario:

- Anti-viral drugs are developed and distributed throughout the United States.
- Development takes 4 years.
- Distribution takes 1 year.
- Remanufacturing and redistribution are required every 10 years because of drug decay over time.

The anti-viral representation differs from the biodetector representation in that the annual representation is either entirely existing measures or entirely anti-viral drugs (Table 4-2) since it is assumed that the entire national supply can be distributed during the year after development and every 10 years thereafter because of the drugs' limited shelf life. The annual cost for implementing this option is $292 million for years 1 through 4; $78 million in years 5, 15, and 25 (for remanufacture and distribution); and zero in other years, as shown in Figure 4-2.

TABLE 4-2 Anti-viral Drug Annual Representation

Year	Weighting of Consequence Case		Amount of Anti-viral Drug Distributed (% of total at beginning of year)
	Existing Measures (%)	Anti-viral Drug (%)	
1	100	0	0
2	100	0	0
3	100	0	0
4	100	0	0
5	100	0	0
6	0	100	100
7	0	100	0
8	0	100	0
…	…	…	…
16	0	0	100
17	0	0	0
…	…	…	…
25	0	100	0

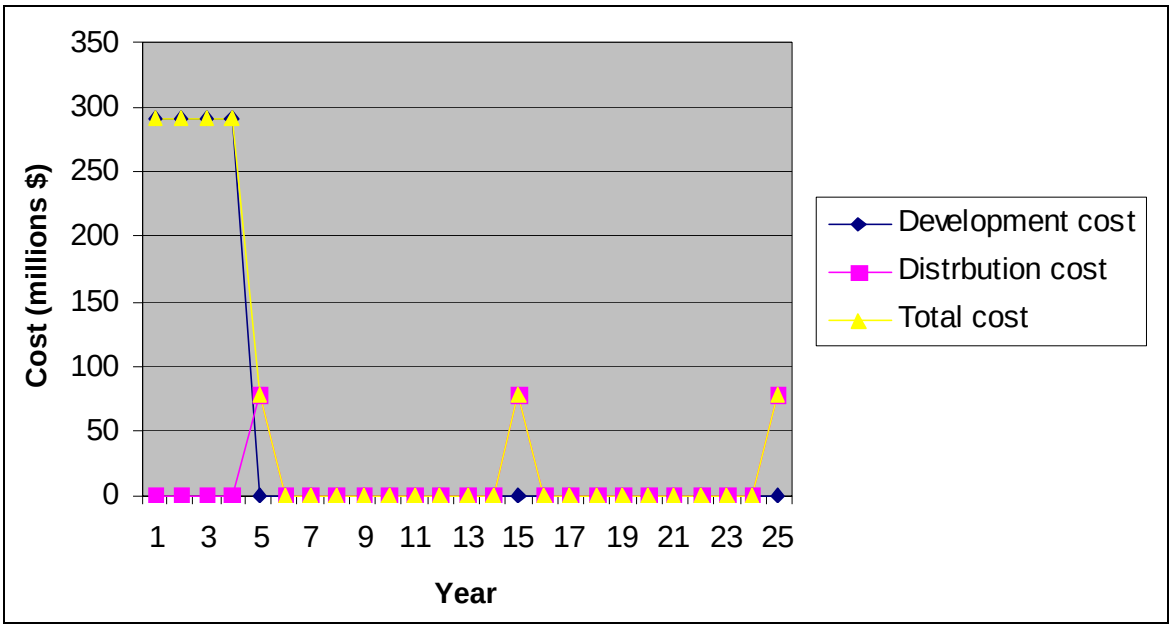

FIGURE 4-2 Annual Cost Distribution for the Anti-viral Alternative

4.4 ALTERNATIVE C: PRETRAINED MEDICAL RESPONDERS

A training program for medical responders and doctors could reduce consequences from those in the existing measures case. The following key assumptions were used for the illustrative pretrained responders scenario:

- 20% of doctors are trained each year in years 1 through 5.
- A new training cycle begins in year 20 for doctors. (This cycle should probably be conducted more often, but this assumption is the one that was made for the case study results that have been completed.)
- 20% of nurses, paramedics, and emergency medical technicians (EMTs) are trained each year on a continuous cycle.
- The self-training media are developed in year 1.
- 70% of health care professionals are trained.
- A combination of live seminars and self-study are used for training.

The annual representation for responders is shown in Table 4-3. The annual cost for implementing this option over years 1–25 ranges from $2.6 million to $4.5 million, as shown in Figure 4-3.

TABLE 4-3 Pretrained Medical Responders Annual Representation

Year	Weighting of Consequence Case: Existing Measures (%)	Weighting of Consequence Case: Pretrained Responders (%)	No. of Doctors Trained during the Year (% of total at end of year)
1	80	20	20
2	60	40	20
3	40	60	20
4	20	80	20
5	0	100	20
6	0	100	0
7	0	100	0
8	0	100	0
…	…	…	…
20	0	0	20
21	0	0	20
…	…	…	…
25	0	100	0

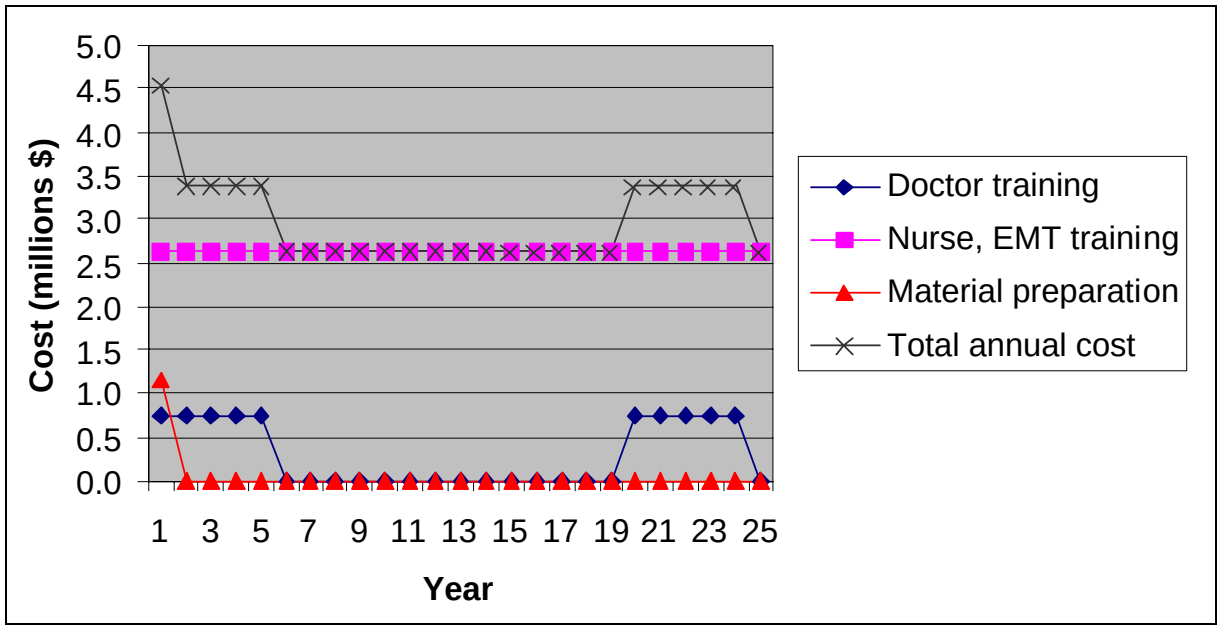

FIGURE 4-3 Annual Cost Distribution for the Pretrained Medical Responder Alternative

5 TIME-DEPENDENT THREAT LEVELS FOR USE IN ILLUSTRATIVE SCENARIOS

As mentioned in Section 3, the likelihood of an incident occurring is linked to the estimated threat levels, which are likely to change over time for a number of reasons. This section briefly describes some differences in incident likelihood between those caused by a natural event and those caused by a terrorist attack. It then presents some alternative representations for incident likelihood for terrorist attacks that can be used to assist in the illustrative comparison of alternative protective measures. The representation of the terrorist threat level presented here is strictly for illustrative purposes and is not linked in any way to actual threat data.

5.1 INCIDENT LIKELIHOOD FOR NATURAL EVENTS AND TERRORIST ATTACKS

An analysis conducted with CIPDSS is intended to assist in risk-informed decision making, which involves the consideration of questions such as these:

- When consequence, vulnerability, and threat information is being incorporated in an overall risk assessment, what are the highest-risk areas?

- What investment strategies can be made that will have the most impact in reducing overall risk?

Although it is not within the CIPDSS scope to estimate the incident likelihoods for various threats, it is implied that researchers who take full advantage of CIPDSS capabilities could incorporate the quantitative interpretation of whatever form of threat information is available to DHS, including information on time-dependent threat levels. Improving one's understanding about alternative courses of action (e.g., investments in protective measures) requires balancing investment costs against reduced risks, which may be measured by metrics such as the number of fatalities, damage costs, level of public confidence, and level of national security.

Obtaining valuable insights from CIPDSS and other methods used to estimate consequences does not require certainty about incident likelihood. However, the ultimate goal of obtaining a better understanding about desirable courses of action (e.g., selecting appropriate protective measures to implement) typically requires some specificity with respect to threat likelihoods.

There are some important differences between the incident likelihood of a terrorist attack on a particular target and the incident likelihood of a catastrophic event caused either by nature or by the failure of man-made structures, machines, or equipment. According to Stern (2003):

> "Terrorist attacks are purposeful, unlike chemical hazards or earthquakes. Moreover, they threaten not only human lives but also political values, interests, and institutions. Government legitimacy is based on the state's monopoly over the use of force and protection of its citizens. Terrorists threaten both of those norms."

A few key observations about the differences follow here:

- The likelihood of a terrorist attack on a particular target is typically related to the consequence level (e.g., public fatalities) of the associated incident. In other words, everything else being equal, terrorists are likely to prefer major events with lots of casualties.

- The likelihood of a terrorist attack on a particular target is typically related to the difficulty in carrying out a successful attack. In other words, everything else being equal, terrorists are likely to prefer less difficult targets.

- The likelihood of a terrorist attack on a particular target is typically related to the level of protective measures implemented. In other words, everything else being equal, terrorists are likely to prefer less difficult targets. If superb preventive measures have been implemented, those measures will probably reduce the likelihood of a terrorist attack on that target.

Everything else being equal, the incident likelihood for a terrorist attack at events that could result in a large number of public fatalities is greater than the incident likelihood for an attack at events that could result in a small number of fatalities. This relationship is contrary to the typical likelihood-of-occurrence relationship that exists for natural events. For example, Table 5-1 shows that the average number of earthquakes per year declines as their magnitudes go up (and the associated likelihood of human losses goes up) and that when an earthquake does occur, the likelihood of occurrence also declines drastically as the magnitude goes up. Similar data are true for most catastrophic events that result from failures of man-made structures, machines, or equipment.

Implementing DHS investment strategies is likely to change the incident likelihood and/or its uncertainty when compared to the situation before they were implemented. For example, the threat of airplane hijacking has presumably been reduced by the introduction of new security procedures, barriers preventing cockpit entry, and wider use of air marshals.

The threat framework used to evaluate protective measures should allow the incident likelihood to be adjusted on the basis of (1) the measures that are assumed to be implemented and (2) the assumptions about improved capabilities of terrorist organizations over time. For those targets protected by the protective measures, the incident likelihood may be decreased. For other similar targets not protected by the measures, the incident likelihood may be increased. The threat framework for evaluating protective measures should also recognize the potential for terrorists, over time, to increase their capability to carry out attacks that have the most desirable consequences from their point of view.

TABLE 5-1 Frequency of Occurrence of Earthquakes as a Function of Magnitude

Qualitative Description	Magnitude (Richter Scale)	Average No. of Occurrences Each Year	Likelihood of Occurrence per Earthquake (% of earthquakes)
Very Minor	2 – 2.9	1,300,000[a]	90
Minor	3 – 3.9	130,000[a]	9
Light	4 – 4.9	13,000[a]	0.9
Moderate	5 – 5.9	1,319[b]	0.09
Strong	6 – 6.9	134[b]	0.009
Major	7 – 7.9	17[b]	0.001
Great	8 and higher	1[c]	0.00007

[a] Estimated.

[b] Based on observations since 1990.

[c] Based on observations since 1900.

Source: USGS (2008)

5.2 REPRESENTATION OF THREAT LEVELS FOR ILLUSTRATIVE ANALYSES OF PROTECTIVE MEASURES

A simplified approach for representing different threat levels has been taken to help demonstrate its importance in evaluating alternative protective measures. Three threat levels were postulated over a 25-year evaluation period:

1. Low (average likelihood of 0.5 incident per 100 years),

2. Middle (i.e., Mid) (average likelihood of 1.5 incidents per 100 years), and

3. High (average likelihood of 5.0 incidents per 100 years).

In addition, for each of the three threat levels listed above, the likelihood was assumed to vary over time in three ways:

1. Constant,

2. Increasing linearly, and

3. Exponentially increasing or S-shaped.

These three time-dependent relationships, used for the hypothetical examples in this report, are not a comprehensive representation of possible threat variation over time. Some threats may decrease over time because of improved countermeasures or emerging new threats. For example, the threat level associated with the use of aircraft as a weapon is arguably lower today than it was in the year 2000.

Figure 5-1 shows the three annual threat curves over a 25-year period for each of the three threat levels. The average value for all three curves is the same, but the linearly increasing and the exponential curves have a lower incident likelihood in the early years and a higher incident likelihood in the later years. The curve shapes have the same form for each of the three threat levels, with different scales for incident likelihood shown on the left in Figure 5-1.

A hypothetical biological incident is used for the illustrative results shown in Section 6. As mentioned in Section 3, terrorists may pursue some types of threat or attack more than others because of the dramatic consequences that could result from such an attack. A biological incident may be one of those preferred threats. It may currently have a low likelihood because of the terrorist's lack of capability, but because of its potential consequences, it may eventually have an higher likelihood. Therefore, most of the examples shown in Section 6 use the exponential threat curve.

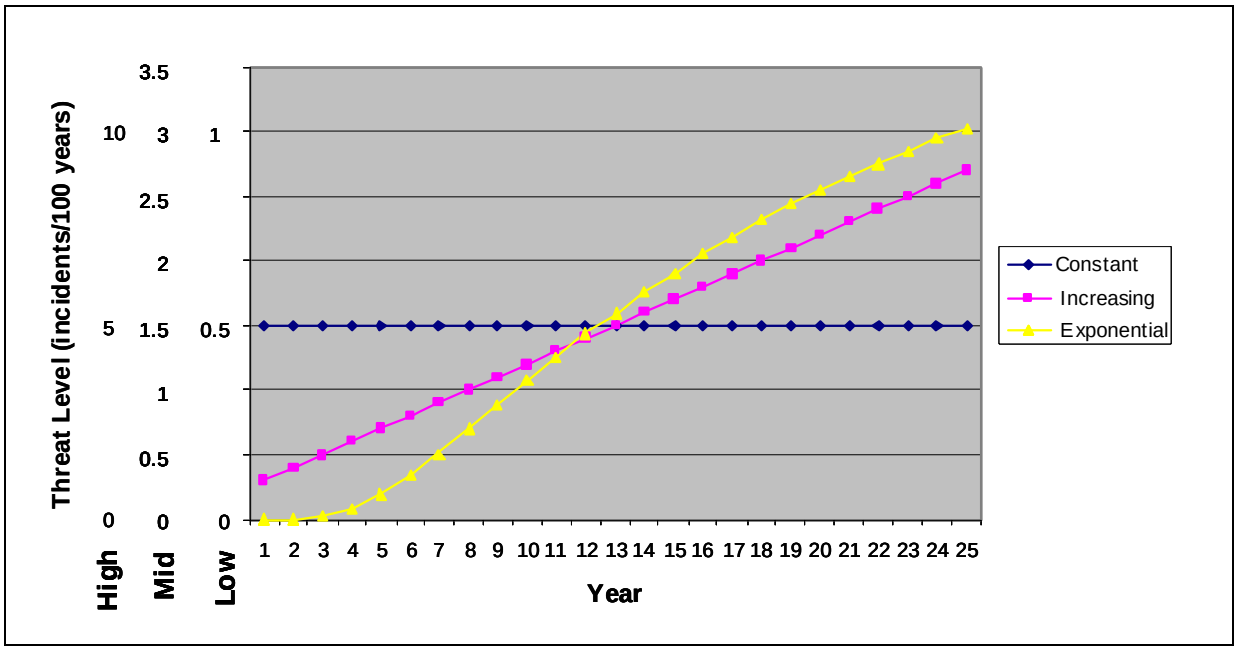

FIGURE 5-1 Alternative Threat Level Representations

6 EVALUATION OF OPTIONS FOR ILLUSTRATIVE BIOLOGICAL THREAT SCENARIOS

A comparison of existing protective measures and the use of detectors, anti-viral drugs, and pretrained responders was conducted for the illustrative biological threat levels outlined in Section 5. All results described in this report use a risk-neutral profile.[3] The process incorporates the following steps:

- Calculate utilities in each year as a combination of utilities of illustrative cases evaluated with the CIPDSS consequence model. (For example, year 1 of the responder alternative is 20% of the responder result plus 80% of the existing measures result; see Table 4-3.) The CIPDSS incident cases accumulate consequences of this incident over a time horizon that continues until a new equilibrium exists and all related impacts have been accumulated.

- Build out the year-by-year table for each alternative for both the readiness and incident scenarios.

- Select a threat profile (Constant, Increasing, or Exponential and Low, Middle, or High).

- Evaluate the year-by-year utility for each alternative. This step is done by estimating the expected consequences in each year by using the incident likelihood and the CIPDSS consequences for an incident case (and the likelihood an incident does not occur and the CIPDSS consequences for the readiness case). Of course, costs and other impacts associated with the protective measures occur in every year independently of whether there is an incident.

- Compare the year-by-year utility for the two alternatives of interest.

- Construct an annual “value of the difference” curve by using the dollar equivalent for the utility.

- Construct a cumulative curve to evaluate how long it takes to break even and thus how long a consistent policy must be supported to show benefits.

The following sections present a few sample comparisons of alternative protective measures for the biological incident case. None of the results shown in these sections involve any discounting of consequences, money, or utilities. The magnitude of the results and the relative

---

3 In general, use of a risk-averse profile causes the breakeven points to shift to the left in Figure 2-3. Roughly speaking, a risk-averse profile will result in a preference to implement the protective measures at lower incident likelihoods than will a risk-neutral profile. The opposite is true for a risk-tolerant profile. The effect of alternative risk profiles is a topic worthy of further consideration, as noted in Section 7.

desirability of the outcomes could change significantly if the assumptions about costs and threats changed.

6.1 PRETRAINED RESPONDERS VERSUS EXISTING MEASURES

Figure 6-1 shows the annual results for the pretrained responders alternative compared to the existing measures alternative (no additional protective measures) for the low exponential threat level (Figure 5-1). The curve shown is the result of subtracting the utility (converted to dollar equivalent) for the existing measures alternative from the pretrained responders alternative. In years when the annual value is greater than zero, the responders alternative is the better performer.

Figure 6-1 shows that the responders alternative is better than the existing measures alternative on an annual basis starting in year 6 and increasingly better in each year after that. In the 25th year, the responders alternative is better than the existing measures alternative by the equivalent of approximately $25 million.

The cumulative results of this comparison are shown in Figure 6-2. The interpretation of this figure gives considerable insight into the potential benefits of one alternative over another. The pretrained responders alternative takes 10 years before it shows positive cumulative benefits over the existing measures alternative for this threat profile. In other words, if continuous support of the pretrained responders policy cannot be maintained for at least 11 years (given the threat

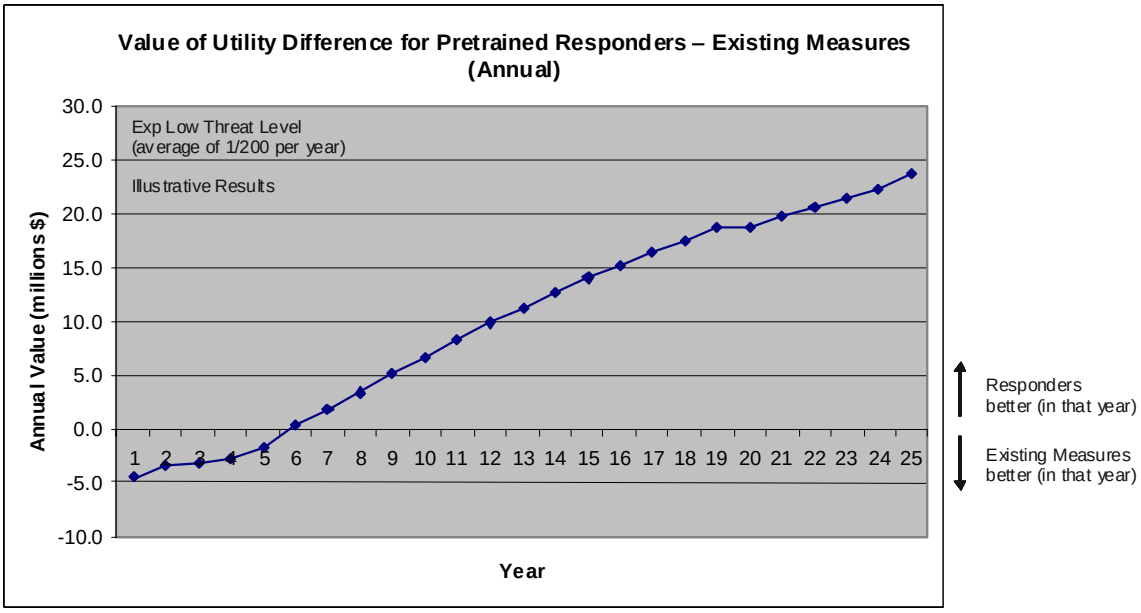

FIGURE 6-1 Annual Comparison of Pretrained Responders Alternative to Existing Measures Alternative

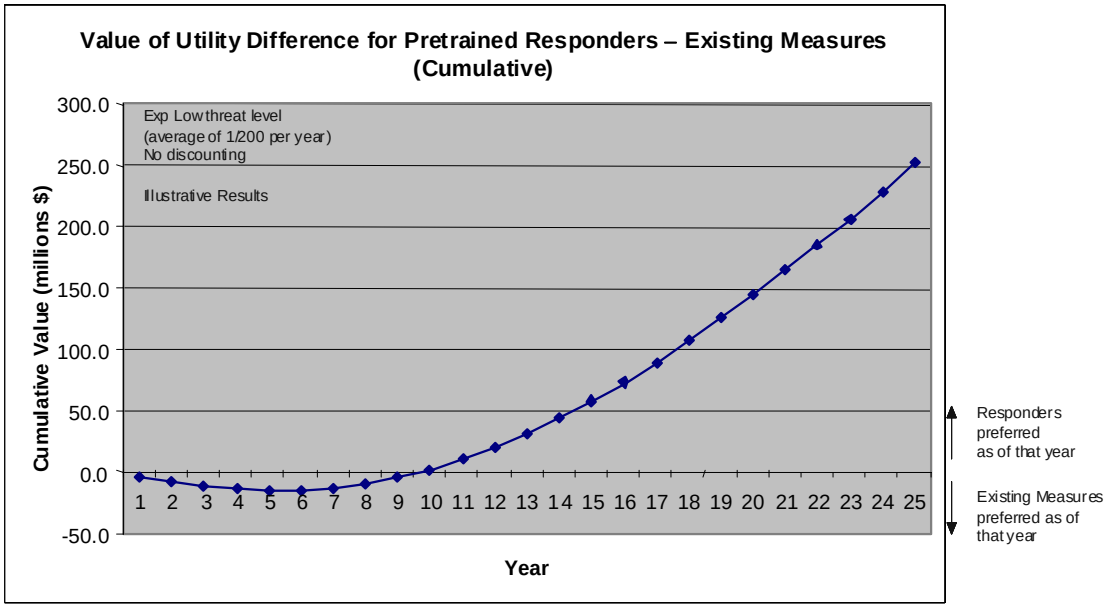

FIGURE 6-2 Cumulative Comparison of Pretrained Responders Alternative to Existing Measures Alternative

level and assumptions about the cost of the policy), then the existing measures alternative is the better choice. The figure shows that if support for the pretrained responders alternative can be maintained for 25 years, given these assumptions, the responders alternative will have accumulated $250 million in benefits.

6.2 DETECTORS VERSUS EXISTING MEASURES

Figure 6-3 shows the annual results for the detectors alternative and existing measures alternative for the middle exponential threat level. The curve shown is the result of subtracting the utility (converted to dollar equivalent) for the existing measures alternative from the detectors alternative. In years when the annual value is greater than zero, the detectors alternative is the better performer.

Figure 6-3 shows that the detectors alternative is better than the existing measures alternative on an annual basis starting in year 13 and is increasingly better each year after that. In the 25th year, the detectors alternative is better than the existing measures alternative by the equivalent of nearly $200 million. The structure in the chart is primarily the result of the initial investment and the replacement of detectors on a 4-year cycle after 8 years of duty.

The cumulative results of this comparison are shown in Figure 6-4. The detectors alternative takes 23 years before it shows more positive cumulative benefits than the existing measures alternative for this threat profile. Figure 4-1 shows that the detectors alternative

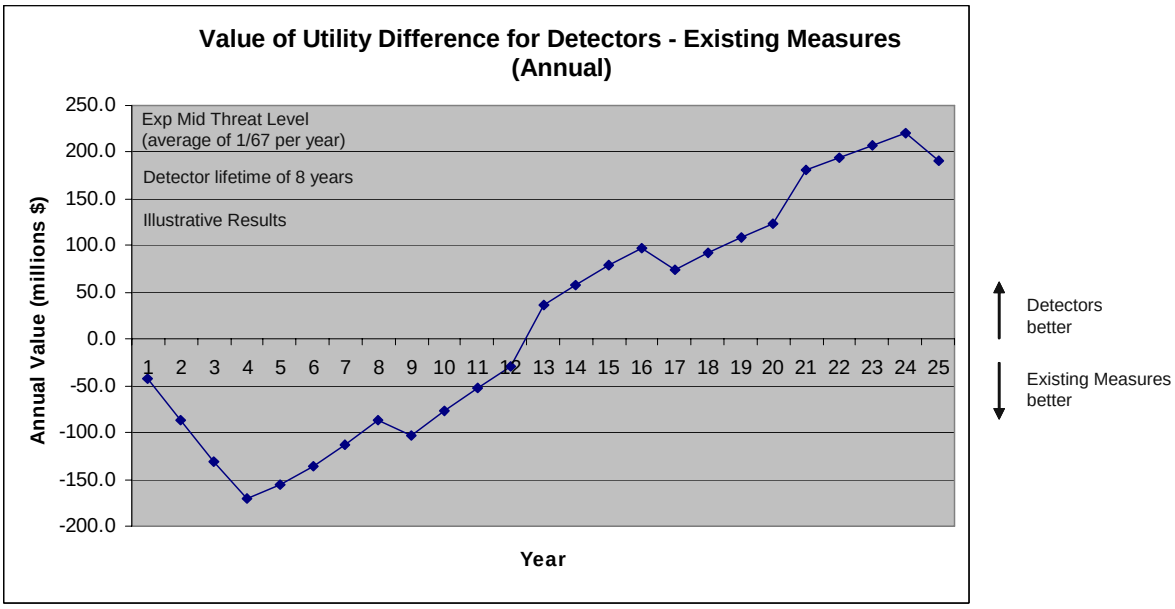

FIGURE 6-3 Annual Comparison of Detectors Alternative to Existing Measures Alternative

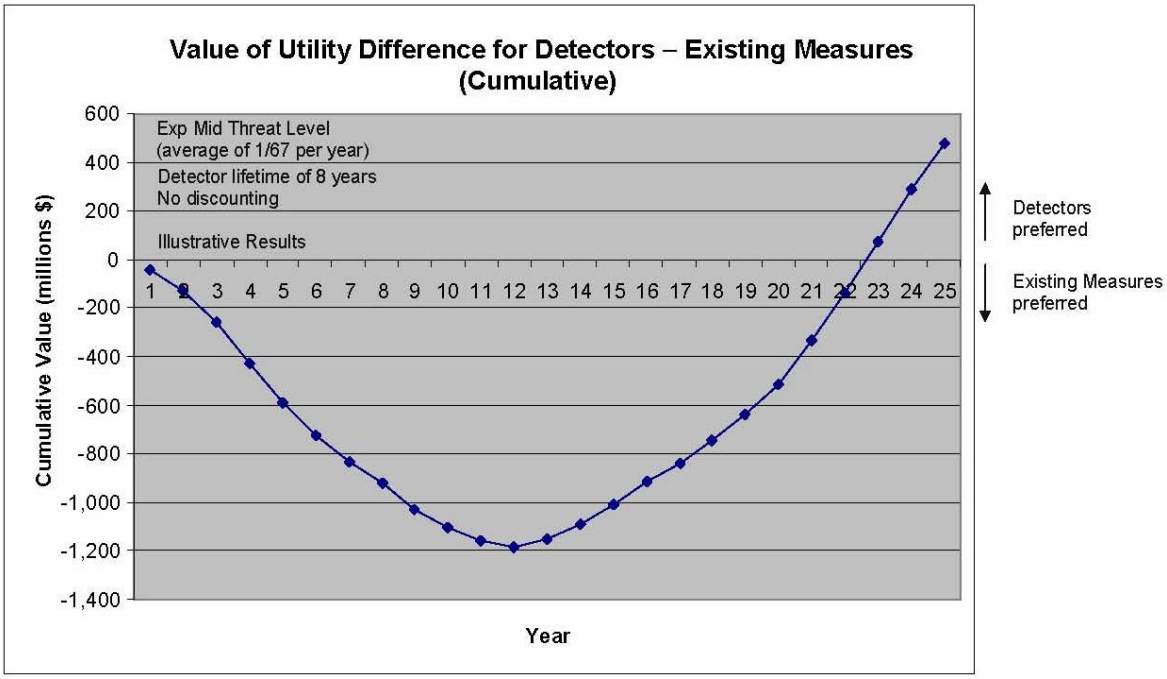

FIGURE 6-4 Cumulative Comparison of Detectors Alternative to Existing Measures Alternative

requires expenditures on the order of $200 million each year. So, if continuous support of the detectors policy cannot be maintained for at least 23 years (given the threat level and cost assumptions), the existing measures alternative is the better choice. Figure 6-4 shows that if support for the detectors alternative can be maintained for 25 years, given these assumptions, it will have accumulated $500 million in benefits. On the other hand, if support for the detectors alternative can be maintained for only 12 years, the existing measures alternative should be supported instead, because it has the equivalent of $1.2 billion in better results.

6.3 RESPONDERS VERSUS DETECTORS

Figure 6-5 shows the annual results for the responders alternative and the detectors alternative for a specific constant threat level (1.85 incidents per 100 years; the reason for this selection will become evident below). The curve shown is the result of subtracting the utility (converted to dollar equivalent) for the detectors alternative from the responders alternative. In years when annual value is greater than zero, the responders alternative is the better performer.

Further, Figure 6-5 shows that the responders alternative is better than the detectors alternative in the early years, when there is heavy spending on detector purchases and installation. Starting in year 5, however, a cyclical pattern results, with the main cause being detector replacement as a result of the 8-year assumed lifetime. After the initial 4 years, the responders alternative is better in the years of replacement detector purchases, and the detectors alternative is better in the other years.

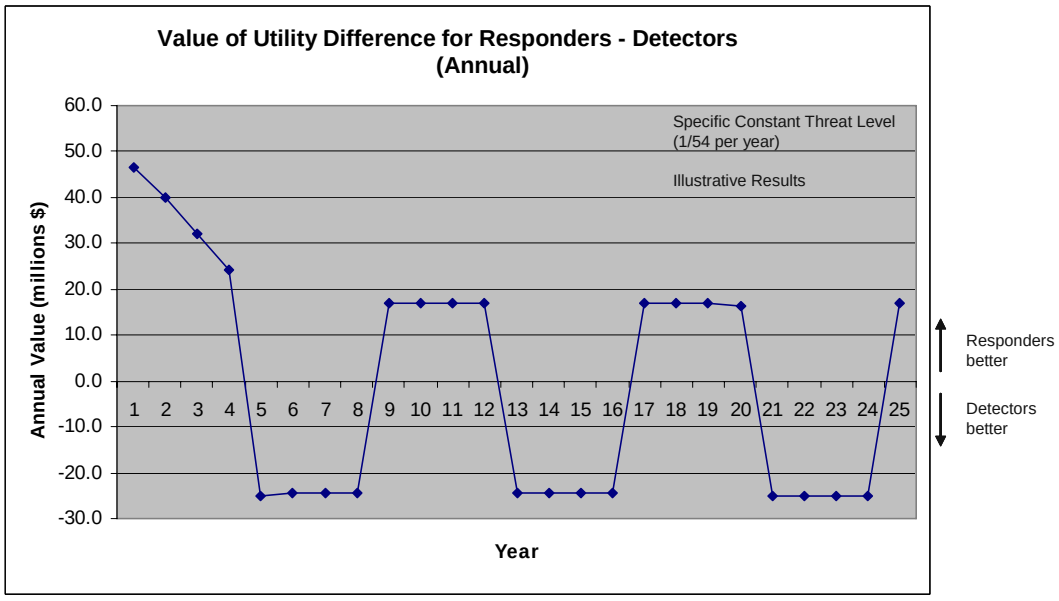

FIGURE 6-5 Annual Comparison of Responders Alternative to Detectors Alternative

The cumulative results of this comparison are shown in Figure 6-6. After 25 years, the two alternatives are exactly even. If the policies to support these two alternatives cannot be extended beyond 23 years, the responders alternative has the better results. If the policies to support these two alternatives could be extended beyond 25 years, an additional analysis would be needed to determine the better selection.

The analysis procedure was used to determine the breakeven threat level, which turned out to be 1.85 incidents per 100 years (23% above the middle threat level and 63% below the high threat level; see Figure 5-1). Thus, in addition to analyzing alternative protective measure scenarios for a specified threat level, this approach can also be used to determine breakeven threat levels.

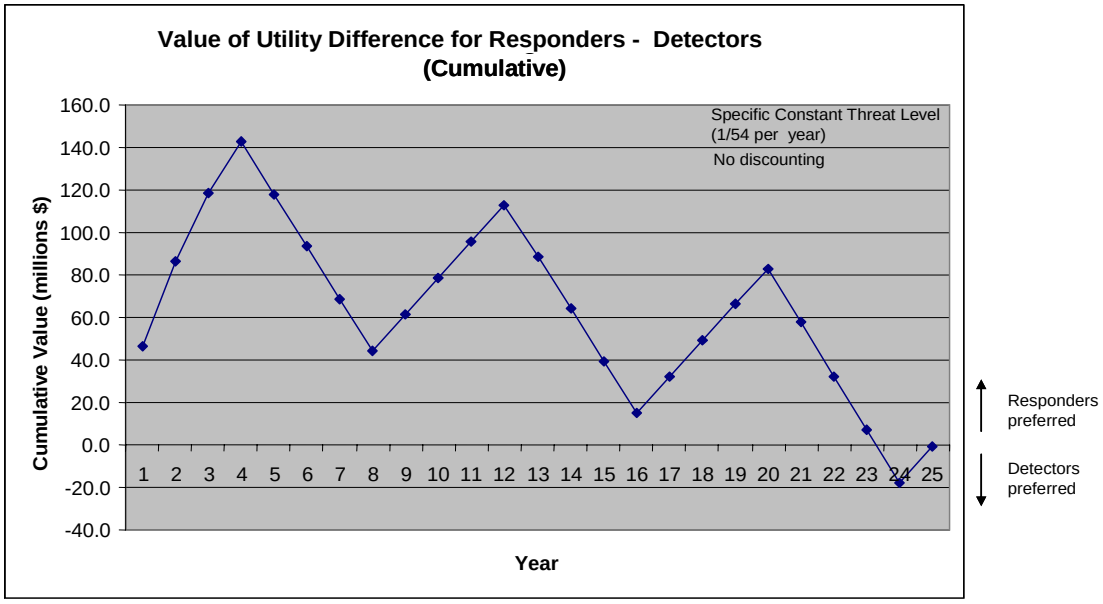

FIGURE 6-6 Cumulative Comparison of Responders Alternative to Detectors Alternative

The breakeven constant threat levels for the three protective measure alternatives considered in the illustrative analysis compared with the existing measures alternative were as follows:

1. Detectors (1.5 incidents per 100 years),

2. Anti-viral drugs (0.84 incident per 100 years), and

3. Pretrained responders (0.12 incident per 100 years).

If the threat level is greater than the breakeven threat level listed, widespread implementation of the applicable protective measure across the United States is better than the existing measures alternative (no additional protective measures).

6.4 RESPONDERS VERSUS DETECTORS AT THE HIGH THREAT LEVEL

Figure 6-7 shows the annual results for the responders alternative compared to the detectors alternative for the high exponential threat level (average of 1/20 per year). The curve shown is the result of subtracting the utility (converted to dollar equivalent) for the detectors alternative from the responders alternative. In years when the annual value is greater than zero, the responders alternative is the better performer.

Figure 6-7 also shows that the responders alternative is better than the detectors alternative in the early years, when there is heavy spending on detector purchases and installation. Starting in year 8, however, the detectors alternative shows an annual advantage that increases significantly over the 25-year period. By the end of this period, the detectors alternative is better than the responders alternative by approximately $900 million per year.

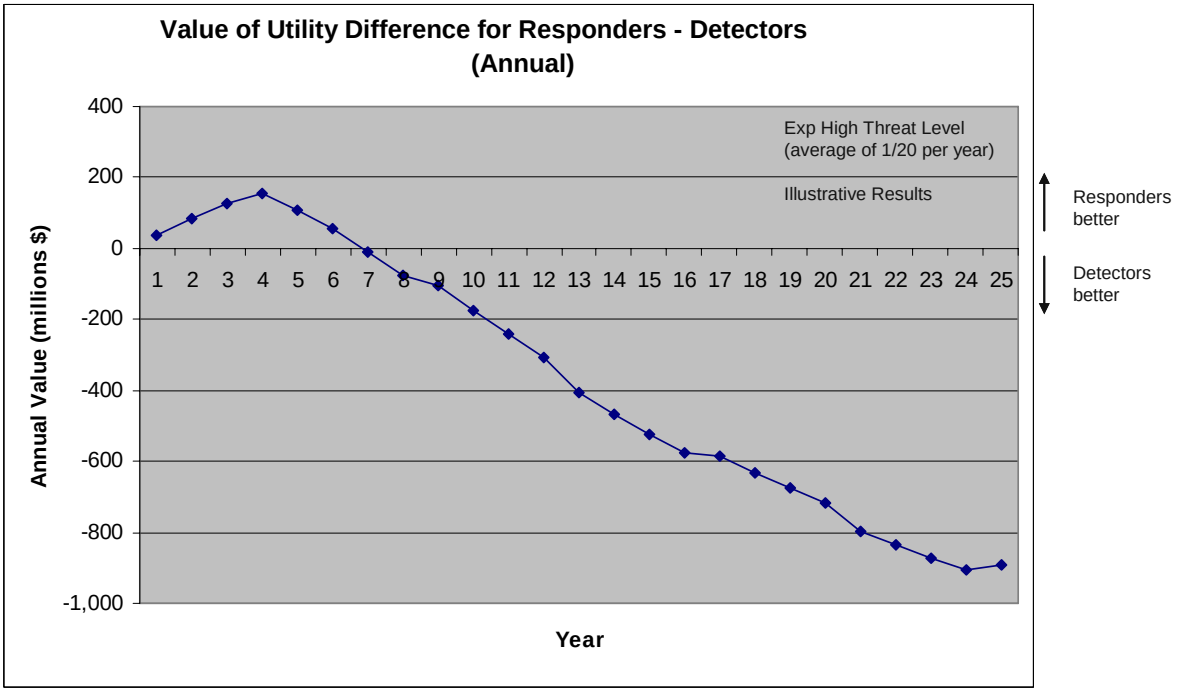

FIGURE 6-7 Annual Comparison of Responders Alternative to Detectors Alternative for the High Threat Level

The cumulative results of this comparison are shown in Figure 6-8. The responders alternative is the better performer through the first 10 years, but then the large benefits of the detectors alternative begin to dominate the results. After 25 years, the detectors alternative has a $9 billion advantage over the responders alternative. Thus, if the detectors alternative policy could receive the continuous support necessary (Figure 4-1), it would have significant long term benefits. If the support for the alternatives could only be maintained for 10 years or less, the responders alternative would be the better choice.

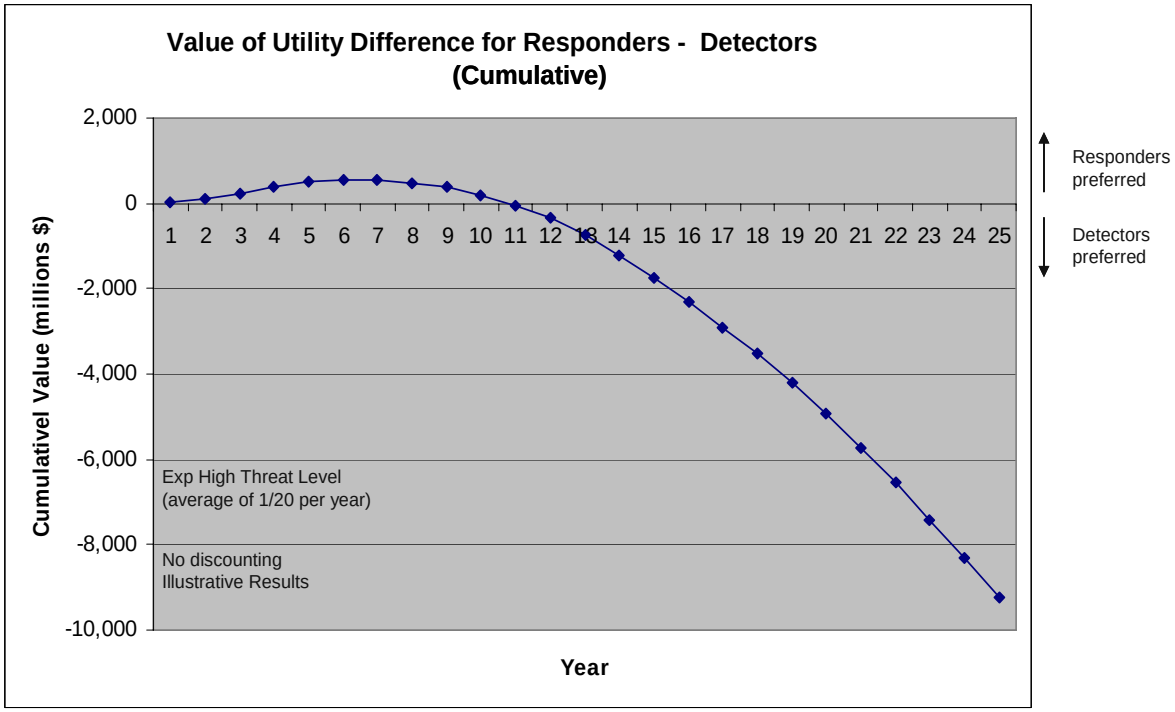

FIGURE 6-8 Cumulative Comparison of Responders Alternative to Detectors Alternative for the High Threat Level

7 ADDITIONAL CONSIDERATIONS IN TIME-DEPENDENT EVALUATION

The factors described in this report are a few of the important time-dependent considerations that may affect the overall merits that result from implementing a particular protective measure. A number of other considerations that have not been examined here are worthy of further consideration. These include:

- Appropriate time frame for evaluation, given the characteristics of protective measures.
- Appropriate time frame for comparison of options.
- Discounting costs and other consequences. (Hypothetical examples given here are intended to illustrate the importance of other time-dependent factors and, in order to avoid further complexity, do not include discounting. Real-world evaluations of alternative protective measures should address discounting explicitly.)
- Introduction of multiple protective measures in a scenario.
- Effects of risk aversion or risk-tolerant behavior on the selection of appropriate alternatives.
- Effective ways to communicate the results in order to compare two options (e.g., years to break even or years of support needed).
- Effective ways to communicate the results for multiple options with multiple regions of dominance.

8 CONCLUSIONS

Through hypothetical examples, time-dependent factors have been shown to affect the relative competitiveness of alternative protective measures. A framework for evaluation that provides some additional capabilities for comparing protective measure options has been described. Although the framework draws extensively on the consequence estimates of the CIPDSS model, it can be applied to any evaluation that uses consequence estimates associated with implementing alternative protective measures.

The characteristics and capabilities of the framework presented here include the following:

- Estimated threat levels that vary with time;
- Implementation times for protective measures;
- Partial capability during implementation;
- Delays in implementation decisions (e.g., funding limits);
- Delays in implementing measures (e.g., physical problems);
- Lifetimes of equipment, training, and medicines;
- Likelihoods of breakeven incidents and the time it takes to break even;
- Annual investment and operational costs and their compatibility with the DHS budget process (versus total present value); and
- Readiness of framework for discounting, if desired.

The results provide insight into how long policies must be supported to show overall benefits. Some examples show that the selection of a different alternative is in order if policy and financial support cannot be maintained for more than 10 years. Some examples show that substantial overall benefits are possible but may not appear until a number of years after the protective measure's initial implementation. The magnitude of the results and the relative desirability of the outcomes could change significantly if the assumptions about costs and threats were different.

9 REFERENCES

DHS (U.S. Department of Homeland Security), 2006, *National Infrastructure Protection Plan*, available at http://www.dhs.gov/xlibrary/assets/NIPP_Plan.pdf.

DHS, 2007, *FY 2007 Homeland Security Grant Program*, available at http://www.dhs.gov/xlibrary/assets/grants_st-local_fy07.pdf.

Keeney, R.L., 1995, "Understanding Life-threatening Risks," *Risk Analysis* 15(6):627–637.

OMB (Office of Management and Budget), 1992, "Guidelines and Discount Rates for Benefit-Cost Analysis of Federal Programs," Circular No. A-94 Revised (Transmittal Memo No. 64), Memorandum from the White House OMB to Heads of Executive Departments and Establishments, Oct. 29, available at http://www.whitehouse.gov/omb/circulars/a094/a094.html.

Stern, J., 2003, "Dreaded Risks and the Control of Biological Weapons," *International Security* 27(3):89–123, winter 2002/2003.

USGS (U.S. Geological Survey), 2008, Earthquake Hazards Program: Earthquake Facts and Statistics, National Earthquake Information Center, available at http://neic.usgs.gov/neis/eqlists/eqstats.html. Accessed May 2008.

Decision and Information Sciences Division
Argonne National Laboratory
9700 South Cass Avenue, Bldg. 900
Argonne, IL 60439-4867

www.anl.gov

UChicago ▶ Argonne LLC

A U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC