2009

Risk-Based Allocation of Resources to Counter Terrorism

Don N. Kleinmuntz
University of Southern California, don.kleinmuntz@usc.edu

Henry H. Willis
hwillis@rand.org

Risk-Based Allocation of Resources to Counter Terrorism

Don N. Kleinmuntz, formerly of University of Southern California
Henry Willis, RAND Corporation
dnk@strata-decision.com
hwillis@rand.org

1. Executive Summary

There is widespread recognition that there is a need to development sophisticated and effective analytic approaches for risk-based allocation of resources to counterterrorism. Implementing these methods requires credible and accurate quantitative assessments of the threats, vulnerability, and consequences of terror attacks, as well as valid assessments of how countermeasures will impact the nature and degree of threat, vulnerability, and consequences. Finally, risk management alternatives must be prioritized in the face of constrained resources – there is simply not sufficient money, time, or human resources available to address all possible countermeasures for all possible threats. Homeland security officials recognize that they need practical, rational, and defensible methods for allocating scarce resources to get maximum protection. In the current (and final year) of this project, there have been two parallel efforts under way.

Robust Portfolio Methods: Experience suggests that efforts to implement these models will often encounter difficulties in obtaining credible inputs. Several difficulties are particularly salient: 1) Quantitative threat assessments expressed as the probability of an attack can be difficult to obtain, although it may be easier to express judgments of relative threat or to rank order threats. 2) Vulnerability assessments of potential targets require expert analyses that can be both expensive and time consuming, particularly when the list of potential targets is long. Resource allocation may have to depend on vulnerability assessments that are incomplete, out of date, or both. 3) Consequence assessments ought to include both direct consequences (fatalities, injuries, damage to property) as well as indirect economic consequences of an attack. While researchers at CREATE and elsewhere have made significant progress in the economic modeling of indirect consequences, it is not uncommon for different estimates to diverge, in some cases across a fairly wide range.

Decision analysis models can become difficult to use or interpret when model parameters are vague and incomplete. Our approach is to identify robust solutions that perform well across a range of plausible parameter values. A traditional way to do this is through sensitivity analysis. A more powerful and compelling alternative is to extend a method called Robust Portfolio Modeling (RPM), previously applied to multi-criteria projects under certainty, to the area of risk-based resource allocation. This is a

"This research was supported by the United States Department of Homeland Security through the Center for Risk and Economic Analysis of Terrorism Events (CREATE) under grant number 2007-ST-061-000001. However, any opinions, findings, and conclusions or recommendations in this document are those of the authors and do not necessarily reflect views of the United States Department of Homeland Security."

computationally intensive approach that relies on a dynamic programming algorithm for computing all non-dominated portfolios of counterterrorism measures, subject to incomplete information about risks and risk management plans (e.g., ordinal threat assessments and/or range-based rather than point estimates other parameters). A basic algorithm for RPM in infrastructure protection has already has been developed and tested for a portfolio of approximately 30 sites. During the current project year, we developed a general approach to assessing threat, vulnerability, and consequences of terror attacks in a form that permits experts to provide initial guidance with minimal time commitment, and in a form that is readily incorporated into the RPM analytical framework. This protocol was tested in an application at California's Governors Office of Homeland Security (OHS), although the nature of that problem did not permit the use of RPM for analysis.

Value of Information Concepts for Selecting Sites for Vulnerability Analysis: A common problem encountered in resource allocation is the scarce nature of resources for performing vulnerability analysis. Starting with only very limited initial information, homeland security agencies must identify where to deploy analytical teams to collect relevant information, assess threats and vulnerabilities, and identify both consequences and the potential to address the vulnerabilities. We developed an approach that would start with the assessment protocol described above to generate high-level assessments of risk for entire sectors, and then evaluate the potential benefits of vulnerability using concepts related to the expected value of perfect information, a well-known decision analysis construct.

The essence of the approach is to analyze a two-stage decision problem, with the initial decision stage being whether to select a site for vulnerability analysis or not, and the second choice being whether or not to expend resources and implement risk management efforts for the site. For vulnerability analysis to be valuable, the decision to implement risk management or not must depend upon the results of the vulnerability analysis – if vulnerability parameters are estimated to be large, then do risk management, while if not, then don't. The model provides a quantitative estimate of the relative benefit of performing vulnerability analysis on a particular site or collection of sites, and this in turn, can be used as an input to an optimization model that determines the optimal set of sites to perform vulnerability analysis upon based on the relative value of the analyses subject to the capacity constraints for vulnerability analysis.

This project is complementary to and provides methodological guidance to applied resource allocation and risk management efforts at CREATE, including potential applied analyses in support of DHS or state/local agencies.

2. Research Accomplishments

2.1. Robust Portfolio Methods for Resource Allocation

Experience suggests that efforts to implement these models will often encounter difficulties in obtaining credible inputs. Several difficulties are particularly salient:

1. Quantitative threat assessments expressed as the probability of an attack can be difficult to obtain, although it may be easier to express judgments of relative threat or to rank order threats.
2. Vulnerability assessments of potential targets require expert analyses that can be both expensive and time consuming, particularly when the list of potential targets is long. Resource allocation may have to depend on vulnerability assessments that are incomplete, out of date, or missing entirely.

3. Consequence assessments ought to include both direct consequences (fatalities, injuries, damage to property) as well as indirect economic consequences of an attack. While researchers at CREATE and elsewhere have made significant progress in the economic modeling of indirect consequences, it is not uncommon for estimates to diverge, in some cases across a fairly wide range.

Decision analysis models can become difficult to use or interpret when model parameters are vague and incomplete. We have developed an approach to identify robust solutions that perform well across a range of plausible parameter values. A traditional way to do this is through sensitivity analysis. A more powerful and compelling alternative is to extend a method called Robust Portfolio Modeling (RPM), previously applied to multi-criteria projects under certainty, to the area of risk-based resource allocation. This computationally intensive approach relies on a dynamic programming algorithm for computing all non-dominated portfolios of counterterrorism measures, subject to incomplete information about risks and risk management plans (e.g., ordinal threat assessments and/or range-based rather than point estimates of model parameters). In previous work, a basic algorithm for RPM in infrastructure protection has been developed and tested with a portfolio of approximately 30 sites.

Most recently, as part of a larger effort to support decisions at California Governor's Office of Homeland Security (OHS), we developed an assessment protocol that permits rapid assessment of key elements of threat, vulnerability, and consequences across many sites in a short period of time. The information collected is useful as an input to robust methods, or for other similar calculations intended to put broad bounds on calculations of risk, and to support identifying infrastructure sites for further analysis.

The heart of the method, when used by a panel of terror risk experts, is to explicitly support vague or imprecise estimates. Threat is assessed using an ordinal rating scale that focuses on the rank ordering of threat probabilities but does not provide precise probability assessments. Vulnerability and consequence assessments involves eliciting ranges (lower and upper bounds), using anchored scales. In the context of the study for OHS, we were focused on assessments of unidentified sites within different infrastructure subsectors:

Threat was defined as the probability of an attack – "Suppose you know an attack would take place in California next year, but the target is unknown. Rate the relative likelihood the attacker would select one or more critical sites in each sector." The scale was from 0 to 10, where 0 was defined as "possible but extremely unlikely" and 10 was defined as "extremely likely."

Vulnerability was defined as the probability an attack succeeds if attempted – "Suppose an attack occurred against a particular site in each sector. Rate the probability that the attack would succeed in causing significant damage, including loss of life and direct or indirect economic losses. Provide both an upper and lower bound." The scale was defined as follows:

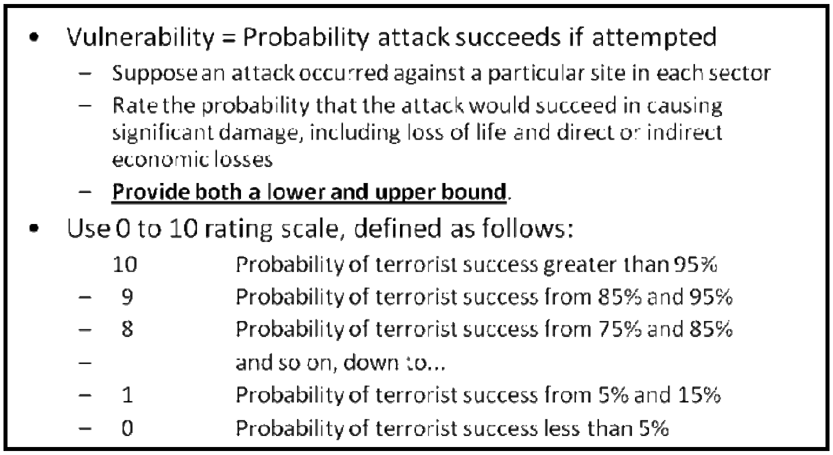

- Vulnerability = Probability attack succeeds if attempted
  - Suppose an attack occurred against a particular site in each sector
  - Rate the probability that the attack would succeed in causing significant damage, including loss of life and direct or indirect economic losses
  - **Provide both a lower and upper bound**.
- Use 0 to 10 rating scale, defined as follows:
  - 10 Probability of terrorist success greater than 95%
  - 9 Probability of terrorist success from 85% and 95%
  - 8 Probability of terrorist success from 75% and 85%
  - and so on, down to...
  - 1 Probability of terrorist success from 5% and 15%
  - 0 Probability of terrorist success less than 5%

Two consequence dimensions were assessed, fatalities and economic loss, with 0 to 7 ratings and both an upper and lower bound assessed:

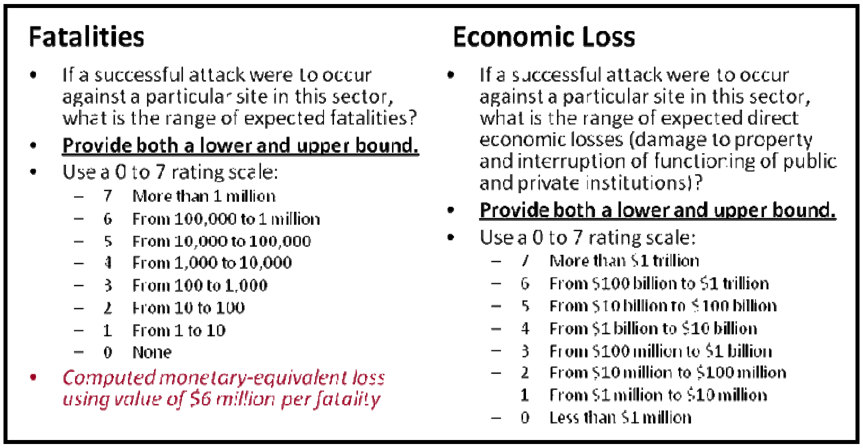

Fatalities

- If a successful attack were to occur against a particular site in this sector, what is the range of expected fatalities?
- **Provide both a lower and upper bound.**
- Use a 0 to 7 rating scale:
  - 7 More than 1 million
  - 6 From 100,000 to 1 million
  - 5 From 10,000 to 100,000
  - 4 From 1,000 to 10,000
  - 3 From 100 to 1,000
  - 2 From 10 to 100
  - 1 From 1 to 10
  - 0 None
- *Computed monetary-equivalent loss using value of $6 million per fatality*

Economic Loss

- If a successful attack were to occur against a particular site in this sector, what is the range of expected direct economic losses (damage to property and interruption of functioning of public and private institutions)?
- **Provide both a lower and upper bound.**
- Use a 0 to 7 rating scale:
  - 7 More than $1 trillion
  - 6 From $100 billion to $1 trillion
  - 5 From $10 billion to $100 billion
  - 4 From $1 billion to $10 billion
  - 3 From $100 million to $1 billion
  - 2 From $10 million to $100 million

  1 From $1 million to $10 million
  - 0 Less than $1 million

Experts from various state and federal agencies felt comfortable in applying the protocol and produced what appeared to be meaningful estimates of ranges of each parameter for varying sites within critical infrastructure subsectors.

2.2. Value of Information Concepts for Selecting Sites for Vulnerability Analysis

A recurrent theme in CREATE's applied work in support of homeland security resource allocation has been that the time and resources available for analysis has been among the scarcest resources available to homeland security agencies. For instance, in a large state like California, time and staffing limitations prevent exhaustive analysis of each and every one of the many hundreds of critical infrastructure sites. This implies the need to do some kind of broad sector-level screening, in order to identify which sector or sectors merits further analysis to determine the most critical sites. Our prior work has been largely *ad hoc*, combining insurance industry databases with a broad and diverse range of other data and subjective inputs. Screening requires a more systematic approach for synthesizing expert opinion and a broad range of data sources in a systematic manner.

We propose using value of information concepts to estimate the value of further investigation of a particular sector, subsector, or individual sites. The idea is to do an initial assessment based on whatever knowledge is at hand (e.g., using the range-based rating protocol described above), and then analysis a two-stage decision problem for each sector: First, should I perform vulnerability analysis on this sector or site (or not). Second, given the results of the vulnerability analysis, should I fund risk reduction to address potential vulnerabilities. Using classic value of information concepts, this amounts to comparing the expected loss reduction benefits associated with two different decision trees: One where we do not know anything beyond our current limited state of knowledge, but must decide whether to invest in protection or mitigation measures, versus another decision tree where we know the results of a vulnerability analysis prior to making the decision about whether to expend resources on protection or mitigation, and the vulnerability analysis results may therefore guide our decision (see below).

The difference between expected losses using one model and the other is the value of doing the analysis prior to the protection/mitigation decision rather than before it. If one is willing to assume that vulnerability analysis reduced the uncertainty about certain model parameters to a point estimate, this corresponds to the Expected Value of Perfect Information (EVPI), an upper bound for the more realistic values that would be achieved given the potential for vulnerability analysis to leave some of the parameters uncertain. If one is willing to assume that vulnerability analysis is no more or no less accurate for different infrastructure sites or sectors, then it is plausible to assume that we can prioritize based on EVPI, and assign our analysts to focus on those sites or sectors with the highest EVPI.

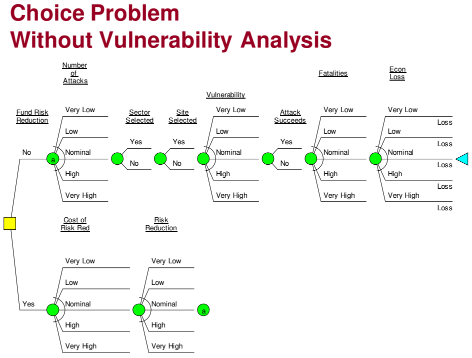
Choice Problem
Without Vulnerability Analysis

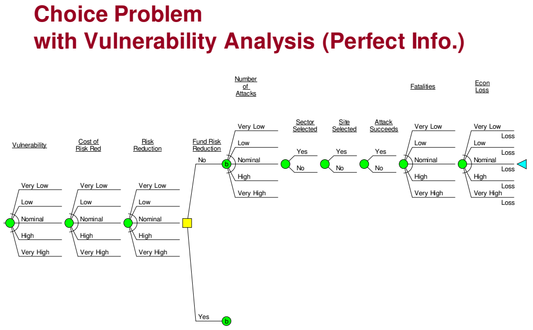
Choice Problem
with Vulnerability Analysis (Perfect Info.)

This method was tested using disguised data from actual California infrastructure sectors. The expected losses from terror attacks were estimated to range up to $6 billion (using a value of $6 million to compute the economic equivalent value of each potential fatality). In a somewhat counterintuitive result, however, we discovered that the value of vulnerability analysis ranged from zero up to a high value of only about $1.1 million. We determined that the value was lower than intuition might suggest because for some infrastructure sectors, the model dictated *always* investing in mitigation/protection independent of whether vulnerability analysis was performed or not. In this instance, it was never a good idea to protect, even if vulnerability analysis suggested very low vulnerability. We hypothesize that this is because the model does not reflect that homeland security officials might be reluctant to authorize protection or mitigation expenditures for particular sites (but not for others) without first performing vulnerability analysis to justify the decisions. When we pruned the "invest in protection or mitigation" branch from the tree computing the expected losses without prior analysis, the value of the analysis was much higher – ranging from zero up to $840 million per site. Under these circumstances, using optimization to assign analysts to the highest value sectors, subject to a limit on the number of sites that can be analyzed ensures that the scarce analyst capacity produces the maximum benefit.

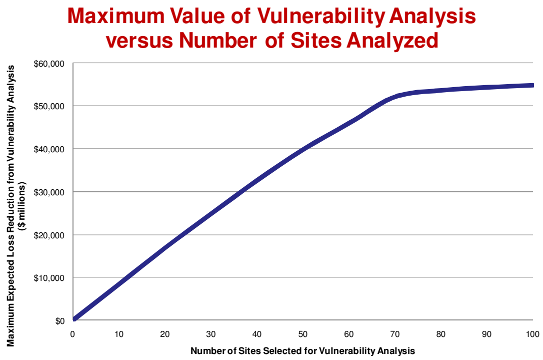

One result still eludes us, however. We had original hoped to be able to use this value of information framework in combination with the Robust Portfolio Method, permitting us to incorporate rank-order probability information into the analysis and considerably reducing the burden placed on the risk assessors. Unfortunately, the non-linear structure of the value of information model makes this extension of the RPM approach very challenging, and has thus far eluded us.

3. Applied Relevance

The entire focus of this stream of work is to develop methods and insights about how to apply rigorous analytical techniques in real-world settings where problems of missing or incomplete information are the rule rather than the exception. The applied work with California OHS has demonstrated that the methods are both relevant and feasible for application. Initial results from both projects were presented to OHS at critical points in their decision making process and did, in fact, provide valuable analytical guidance to their decision making.

4. Collaborative Projects

All work on these projects would not have been possible without the guidance and assistance of the California Governor's Office of Homeland Security, who helped to guide the definition of the research questions, and who provided assistance in risk assessment and in evaluating the results.

5. Research Products

Research Products (Please detail below)		#
5a	# of non-peer reviewed publications and reports	1
5a	# of scholarly journal citations of published reports	5
5b	# of scholarly presentations (conferences, workshops, seminars)	16
5b	# of outreach presentations (non-technical groups, general public)	7

5.1. Presentations - Conferences

1. Kleinmuntz, D., "Robust Portfolio Methods for Counterterrorism Resource Allocation: Prioritizing Vulnerability Analyses for Critical Infrastructure Sectors," *INFORMS Annual Meeting*, San Diego, CA, October 13, 2009
2. Kleinmuntz, D., "Prioritizing Terrorism Vulnerability Analyses for Critical Infrastructure Sectors," *EURO European Operational Research Conference*, Bonn, Germany, July 6, 2009
3. Kleinmuntz, D., "Resource Allocation Models for Terrorism Risk Management," *Finnish Operations Research Society Meeting*, Helsinki, Finland, November 13, 2008

5.2. Presentations - Outreach

1. Kleinmuntz, D., "Critical Infrastructure Site Prioritization and Resource Allocation," *Executive Program*, CREATE, USC, Los Angeles, CA, July 23, 2009
2. Kleinmuntz, D., "Robust Portfolio Methods for Risk-Based Counterterrorism Resource Allocation," *DHS University Summit*, Washington, DC, March 17, 2009

6. Education and Outreach Products

Education and Outreach Initiatives (Please detail below)	**#**
# of contacts with DHS, other Federal agencies, or State/Local (committees)	15 - 20

Multiple contacts with California Governor's Office of Homeland Security Infrastructure Division over a period of three years, including presentations, research meetings, and project-related discussions.

Two-day contact with DHS Office of Risk Assessment to assist in evaluation of DHS risk assessment methodologies.

Several meetings with Port of Los Angeles / Port of Long Beach security officials to discuss risk assessment and risk management approaches.