This product is part of the RAND Corporation monograph series. RAND monographs present major research findings that address the challenges facing the public and private sectors. All RAND monographs undergo rigorous peer review to ensure high standards for research quality and objectivity.

Estimating Terrorism Risk

Henry H. Willis, Andrew R. Morral,
Terrence K. Kelly, Jamison Jo Medby

RAND CENTER FOR TERRORISM RISK MANAGEMENT POLICY

The research described in this report results from the RAND Corporation's continuing program of self-initiated research. Support for such research is provided, in part, by donors and by the independent research and development provisions of RAND's contracts for the operation of its U.S. Department of Defense federally funded research and development centers.

Library of Congress Cataloging-in-Publication Data

Estimating terrorism risk / Henry H. Willis ... [et al.].
p. cm.
"MG-388."
Includes bibliographical references.
ISBN 0-8330-3834-6 (pbk. : alk. paper)
1. Terrorism—United States—Prevention. 2. Terrorism—Risk assessment—United States. 3. Federal aid to terrorism prevention—United States—Planning. I. Willis, Henry H. II. Rand Corporation.

HV6432.E78 2005
363.32—dc22

2005024093

Cover photo: Getty Images

The RAND Corporation is a nonprofit research organization providing objective analysis and effective solutions that address the challenges facing the public and private sectors around the world. RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

Published 2005 by the RAND Corporation
1776 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138
1200 South Hayes Street, Arlington, VA 22202-5050
201 North Craig Street, Suite 202, Pittsburgh, PA 15213-1516
RAND URL: http://www.rand.org/
To order RAND documents or to obtain additional information, contact
Distribution Services: Telephone: (310) 451-7002;
Fax: (310) 451-6915; Email: order@rand.org

Summary

The Department of Homeland Security (DHS) is responsible for protecting the United States from terrorism through prevention, preparedness, and response. In part, this goal is achieved through allocation of resources to states and urban areas. The Urban Areas Security Initiative (UASI) is a DHS grant program designed to enhance security and overall preparedness by addressing unique planning, equipment, training, and exercise needs of large urban areas (DHS, 2004). Although there appears to be agreement among many stakeholders that these allocations should reflect the magnitude of risks to which different areas are exposed, no consensus has emerged on how this might be accomplished. Indeed, the UASI grant program has frequently been criticized for inadequately calculating risk and therefore for failing to distribute resources in proportion to urban areas' shares of total terrorism risk.

Ultimately, efficient allocation of homeland security resources would be determined based upon assessment of the cost effectiveness of alternative risk-reduction opportunities. After potentially first addressing obvious and easily mitigated risks, this requires understanding the cost effectiveness of different types and amounts of investment. Neither the methods nor the data are available to answer questions about the effectiveness of available risk-reduction alternatives or to determine reasonable minimum standards for community preparedness. Until these questions are answered, allocating homeland security resources based on risk is the next best approach since areas at higher risk are likely to have more and larger opportunities

for risk reduction than areas at lower risk. That is, resources would be allocated roughly proportionally to the distribution of risk across areas receiving funding.

This monograph offers a method for constructing an estimate of city risk shares, designed to perform well across a wide range of threat scenarios and risk types. It also proposes and demonstrates a framework for comparing the performance of alternative risk estimates given uncertainty in measuring the elements of risk.

Components of Risk

Terrorism risk can be viewed as having three components: the *threat* to a target, the target's *vulnerability* to the threat, and the *consequences* should the target be successfully attacked. People and organizations represent threats when they have both the intent and capability to damage a target. The *threats* to a target can be measured as the probability that a specific target is attacked in a specific way during a specified period. Thus, a threat might be measured as the annual probability that a city's football stadium will be subject to attack with a radiological weapon.

Vulnerability can be measured as the probability that damage occurs, given a threat. Damages could be fatalities, injuries, property damage, or other consequences; each would have its own vulnerability assessment. *Consequences* are the magnitude and type of damage resulting, given a successful terrorist attack. Risk is a function of all three components: threat, vulnerability, and consequences. These constructs can be used to measure risk consistently in terms of expected annual consequences. More detailed definitions of vulnerability, threat, and risk and discussions of measures for each are presented in this monograph.

Uncertainty and Value Judgments in Terrorism Risk Assessment

There are two important sources of uncertainty in estimating terrorism risk. The first includes variability and error in estimates of threats, vulnerabilities, and consequences. The second involves how we should value different types of consequences. Part of an informed discussion of homeland security policy rests on an understandable and transparent means of accounting for uncertainties in estimates and the consequences of using alternative values.

When facing uncertainty about estimates and values, policy analysis often relies on best estimates, even when they have a low probability of being correct—and a high probability of being wrong. While this allows us to generate a very precise estimate of risk, in the end, if the estimates poorly represent what actually happens in real life, the precision is misplaced. So, rather than seek an optimal method for estimating risk, we seek a method that leads us to make the least egregious errors in decisionmaking across the range of possible scenarios that might develop in the future. Following methods of adaptive planning under deep uncertainty (Davis 1994, 2002; Lempert, Popper, and Bankes, 2003), we seek a method for estimating risk that is robust because it has the lowest expected error when evaluated against a wide range of possible futures.

One approach for developing an estimate with these properties would be first to define multiple sets of threats, vulnerabilities, and consequence measures and use them as the basis for constructing a single risk estimate. Then using these multiple estimates, develop a single description of how risk is distributed that balances across multiple perspectives of terrorism risk.

Generating multiple risk estimates can provide plausible bounds on the magnitude of terrorism risk estimates and how different stakeholders may be affected based on where they live or what outcomes they value most. From several estimates, one can ask how low or high terrorism risk may be in a specific city, what a best estimate of risk is given the range of estimates available, and how answers to these questions differ when considering different types of outcomes. The challenge to analysts defining a single picture of how risk is distributed is to do so without losing significant information. This requires specifying how to deal with the challenges of aggregation given both inherent uncertainties and value choices.

Uncertainty in terrorism risk estimates suggests the need to devise means of hedging our homeland security policies against a range of distributions of risk that are plausible given what we know about uncertainties in our risk estimation procedures. So, rather than seek an optimal method for estimating risk, we seek a method that leads us to make the least egregious errors in decisionmaking across the range of possible scenarios that might develop in the future. This presents a problem comparable to that of forecasting economic trends using multiple estimates or models discussed by Clemen (1989). This literature highlights two objectives to consider when combining estimates: 1) use information contained in the multiple estimates to improve forecasting accuracy and 2) make note of and retain the important distinctions that individual estimates represent.

Addressing multiple values or objectives in terrorism risk estimates differs from combining forecasts. While the goal of combining forecasts is to develop an accurate estimate, the goal of considering multiple objectives is to reflect appropriately the range of values held by stakeholders. Literature on multiobjective decisionmaking provides several approaches for addressing the fact that terrorism risk can be expressed in multiple outcomes. The commonality across these methods is the need to reflect transparently a range of values for multiple objectives in the decisionmaking process.

Simple Versus Complex Risk Indicators

Despite the many sources of uncertainty surrounding terrorism risk, estimating this risk is necessary for informed distribution of homeland security resources. Approaches that have been used in policy analysis for estimating terrorism risk are bounded by two generic categories: simple risk indicators and event-based models. Each approach reflects the components of terrorism risk (i.e., threat, vulnerability, and consequences) and their uncertainties in different ways. As examples of simple indicators, we describe how population and density-weighted population have been used as estimates of terrorism risk. As an example of event-based models, we describe the Risk Management Solutions (RMS) Terrorism Risk Model. These two examples allow for comparisons that illustrate the strengths and weaknesses of each approach.

There is a logical link between population-based indicators and terrorism risk. An argument can be made that consequences are correlated with population and threats are correlated with population density. There are practical benefits for using simple risk indicators such as those based upon population. In general, the metrics for measuring these indicators are well understood and measurable, and data is widely available. The main limitation of these simple indicators is that they do not fully reflect the interactions of threat, vulnerability, and consequences. As a result, there is little consensus and no validated framework for deciding how to use several simple indicators to create a single risk estimator.

Event-based models are built upon relatively detailed analysis of consequences from specific attack scenarios. These models include sensitivity analysis for important parameters that affect consequences. They may include components to model multiple types of events and multiple targets. They may also include modules that translate expert judgments of likelihood or consequences. The strength of event-based models lie in the greater fidelity they enable in analysis. The weakness is that, to obtain this detail, analysts must estimate many uncertain parameters. One example of an event-based model is the RMS Terrorism Risk Model. The RMS model, discussed in more detail later in this monograph, was developed as a tool for the insurance and reinsurance industries to assess risks of macroterrorism.[1]

---

[1] RMS defines *macroterrorism* as attacks capable of causing (1) economic losses in excess of $1 billion, or (2) more than 100 fatalities or 500 injuries, or (3) massively symbolic damage.

Estimating Terrorism Risk

To demonstrate an approach for estimating terrorism risk to inform resource allocations, we calculated a single estimate of U.S. cities' risk shares based on multiple perspectives of terrorism risk obtained from the RMS Terrorism Risk Model. By considering three perspectives on threat (the RMS standard and enhanced and reduced threat outlooks), the RMS results provide three estimates of terrorism risk for the urban areas that received UASI funding in fiscal year 2004. Using these three sets of expected fatalities, we calculated an aggregated risk estimator by minimizing the sum of the squared underestimation error across all urban areas and risk estimates. The assumptions and limitations inherent to this approach are presented in the monograph.

Strengths and Weaknesses of Different Risk Indicators

We compared the proposed aggregated risk estimator to simple population-based indicators, looking at both how the distribution of risk changes for each and the propensity for each to underestimate a city's risk share given the uncertainty that surrounds terrorism risk.

The aggregated risk estimator concentrated most of the expected terrorism losses in relatively few cities compared to population-based indicators. In addition, the aggregated estimator resulted in the lowest underestimates of risk aggregated across all urban areas. Density-weighted population performed better than population alone and was, in fact, quite a bit closer in performance to that of the aggregated indicator than to that of the simple population indicator.

Because the density-weighted population indicator performs well and is easier to derive than the event-based indicator, it might be of utility for some purposes, e.g., in risk-based allocation of resources for strategic purposes over long time intervals, during which relative risk across urban areas is not expected to change much.

Density-weighted population, however, does not allow decisionmakers to see how changes in threat or vulnerability information affect terrorism risk. For example, when making operational resource allocations or evaluating the effectiveness of preparedness programs, decisionmakers need to understand how specific countermeasures reduce or change the profile of terrorism risk. Similarly, a crude indicator like density-weighted population would offer no guidance about how city risk estimates might change with, for instance, new intelligence about terrorist targeting or capabilities of using weapons of mass destruction (WMD). For these purposes, more detailed event-based models of terrorism risk are essential.

In this study, a single event-based estimate was shown to be robust across uncertainties about the likelihood of WMD attacks, uncertainties about which consequences ought to be prioritized in considerations of city risk, and uncertainties about the expected magnitude of risks each city might face. For example, an important observation is that the risk profile of the urban areas examined did not change significantly with the variability of threats from weapons of mass destruction. This is clearly a function of the models used in this study and how they were parameterized. While the primary focus of this study was not to estimate precisely terrorism risk in the United States, this observation raises questions such as whether risk is a characteristic of a region's infrastructure or population that is relatively stable across different threats. If so, this would be an important observation when it comes to policy and resource decisions.

Recommendations

Our framework for defining terrorism risk and the analysis we present here lead us to five recommendations for improving the allocation of homeland security resources:

1. DHS should consistently define terrorism risk in terms of expected annual consequences. Calculating expected annual consequences requires accounting for threat, vulnerability, and consequences. Defining terrorism risk in these terms facilitates the incorporation of risk reduction as the goal of homeland security programs.

2. DHS should seek robust risk estimators that account for uncertainty about terrorism risk and variance in citizen values. Given the tremendous uncertainties surrounding terrorism risk assessment, it is prudent to plan for the range of plausible futures that may play out. Several approaches are available for generating estimates of city risk shares that offer robust characterizations of risk across multiple uncertainties and perspectives on relative values of different consequences. Our approach to this problem ensures that underestimation error is minimized.

3. DHS should develop event-based models of terrorism risk, like that demonstrated in this monograph. Measuring and tracking levels of terrorism risk is an important component of homeland security policy. These data provide insight into how current programs are reducing risk and when and where new terrorist threats may be emerging. Only event-based models of terrorism risk provide insight into how changes in assumptions or actual levels of threat, vulnerability, and consequences affect risk levels.

4. Until reliable event-based models are constructed, density-weighted population should be preferred over population as a simple risk indicator. Density-weighted population is reasonably correlated with the distribution of terrorism risk across the United States, as estimated by event-based models like the RMS Terrorism Risk Model. To support strategic policy decisions when the effects of new countermeasures or recent intelligence are not in question, density-weighted population is a useful indicator in lieu of event-based models. In contrast, our results suggest that population offers a remarkably weak indicator of risk, not much superior to estimating risk shares at random.

5. DHS should fund research to bridge the gap between terrorism risk assessment and resource allocation policies that are cost effective. We do not here seek to understand how UASI allocation amounts may reduce risk. Until that relationship is understood, resource allocation decisions will not be optimized for reducing casualties and property loss. To these ends, DHS should evaluate the performance of the formula used to assign UASI grants using the approach presented in this study.