This article was published in a Special Issue of *Reliability Engineering and System Safety*, Vol. 86, no. 2, November 2004.

Confronting the Risks of Terrorism: Making the Right Decisions

Prepared by
A Special Study Group on Combating Terrorism,
B. John Garrick, Chair

2004

Confronting the Risks of Terrorism: Making the Right Decisions[1]

SPECIAL STUDY GROUP ON COMBATING TERRORISM

B. John Garrick, Chair, Independent Consultant, 221 Crescent Bay Drive, Laguna Beach, CA 92651, USA[2]
James E. Hall, Hall and Associates
Max Kilger, Symmetrical Resources
John C. McDonald, MBX, Inc.
Tara O'Toole, University of Pittsburgh Medical Center
Peter S. Probst, Independent Consultant
Elizabeth Rindskopf Parker, McGeorge School of Law, University of the Pacific
Robert Rosenthal, Booz Allen Hamilton
Alvin W. Trivelpiece, Oak Ridge National Laboratory (director emeritus)
Lee A. Van Arsdale, Unconventional Solutions, Inc.
Edwin L. Zebroski, Electric Power Research Institute (retired)

STAFF TO THE STUDY GROUP

Raphael Perl, U.S. Congressional Research Service

CONSULTANTS

Stanley Kaplan, Bayesian Systems, Inc.
John W. Stetkar, Independent Consultant

[1] The study resulting in this report was initiated by The National Academy of Engineering of the United States and was conducted by a study group with a wide range of expertise in issues relating to risk and terrorism. A decision was ultimately made not to publish the report as an Academy document but rather as a personal statement by the individual authors. They, and not the Academy, are solely responsible for its content.

[2] Corresponding author. E-mail: bjgarrick@aol.com.

TABLE OF CONTENTS

TABLES AND FIGURES

PREFACE

This report was prepared during a time of unprecedented change in the security of the United States. New laws have been enacted, departments and agencies have been realigned, tens of thousands of new security personnel have been hired, and public awareness of the threat of terrorism has penetrated every community.

The purpose of this report is to suggest a methodology for assessing the risk of catastrophic terrorist attacks, i.e., high consequence attacks that would result in significant loss of life and/or economic damage. While there are numerous examples of quantitative risk assessment and management methods being applied throughout government and the private sector, there has been so far only limited application of the discipline of *Quantitative Risk Assessment*, which emphasizes the quantification of uncertainties based on the available evidence. *The ability to quantify uncertainties is the key to understanding those rare, but critical terrorist attacks that may have catastrophic consequences.* The purpose of this report is to present such a methodology.

The broad context to which the methodology is applied for analyzing the risk of terrorism is important. The study group realizes that this is not the final word on how to analyze the risk of a terrorist attack. Rather it is "work in progress." For this reason care is taken: (1) to emphasize the category of terrorist threats having the potential for catastrophic consequences, (2) to suggest a methodology that is general enough to be applied to a variety of terrorist initiated events, and (3) to provide a methodology that has sufficient analytical muscle to dig much deeper than the usual qualitative methods.

The report also emphasizes the need for proper use of all available information in conducting a quantitative risk assessment (QRA) and the need for appropriate organizational responses to create a successful terrorism risk management program. This report is based on well established methods of risk assessment as used in fields such as nuclear power, the chemical and petroleum industry, and more recently, the space program. Applying such methods provides a much-improved basis for making the "right decisions" about how to combat terrorist attacks that could have catastrophic consequences. The primary audience for this report is policy makers and decision makers in government and industry, but the report also reaches out to practitioners.

A major problem in combating terrorism is ensuring that the public and private sector invest resources rationally in ways that actually reduce the threat and vulnerabilities in all segments of society. A logical approach then would be to take timely, investment-wise steps that not only reduce the threat of terrorist attacks occurring, but also lessens the vulnerabilities to attacks that do occur. The implementation of such an effective strategy will depend on leaders in government and industry understanding the risk quantitatively, and out of this understanding, making the right investments and interventions.

The requirements for making and executing good decisions include: (1) a clear understanding of the nature and characteristics of the terrorist threat; (2) a methodology that systematically and quantitatively exposes and assesses the terrorist threats (anticipated attack

scenarios) and the vulnerabilities to them; (3) an information base relevant to the issues and decision options being considered; and (4) organizational structures and relationships that facilitate both understanding and implementation of the decisions made.

In support of these objectives this report gives an example application of managing the risks of terrorism using an electrical power grid as a case study. The example involves a combined physical and cyberattack on a regional electrical power grid. The vulnerabilities of the grid are systematically exposed, and corrective actions for reducing the risk are identified. The example is purposefully simplified, to communicate understanding of the basic ideas.

Information and supporting evidence are critical to quantifying the risks of terrorist attacks. The information useful for combating terrorism is often fragmented, limited in scope, and not systematically linked or integrated. Moreover, prior to 9/11, the timely sharing of information was not considered an issue because we did not anticipate a serious terrorist attack. Now we understand, better, that sharing information between government agencies, between government and industry, and with the public is crucial to our security.

An important issue addressed in this report is the challenge of promoting organizational relationships and institutional mechanisms for reducing terrorist threats that are hidden and difficult to detect. Nevertheless, there are many examples of how organizations can, and have dealt with threats especially with respect to technological issues. In the 1940s, resources were mobilized for the Manhattan Project and the building of the atomic bomb. In the 1950s, the private sector joined the federal government in responding to the challenge of Sputnik, and we landed a man on the moon by the end of the next decade. In the 1970s, in response to the challenges of trade in high technologies developed around the world, U.S. companies established partnerships with universities and national laboratories. In the l980s, when U.S. companies appeared to be losing the lead in the integrated circuit industry, federal action was taken to break down antitrust and other legal barriers to cooperation among key industry firms, and SEMATECH, a public-private partnership, was established. One of the challenges facing the nation today is to create organizational relationships that mobilize our engineering, scientific, and technological communities, under government leadership, to counter the terrorist threat.

The study group has focused on methodologies for assessing the risk of catastrophic terrorist attacks. The study group preparing this report included 11 individuals from 10 different segments of our society (see biographies in Appendix C). The study group had staff support as well as two outside consultants. The conclusions are based on unclassified briefings from government, academic, and industry experts (see Appendix D) and from exchanges of views based on study group members' experiences and expertise.

B. John Garrick, *Chair*
Study Group on Combating Terrorism
221 Crescent Bay Drive
Laguna Beach, CA 92651, USA
bjgarrick@aol.com

EXECUTIVE SUMMARY

The events of September 11, 2001, demonstrated to the nation and the world the vulnerability of the United States to catastrophic terrorist attacks. This attack has affected the life-style of every American in terms of work, travel, planning, and leisure activities. We have had to adopt a new consciousness about threats and vulnerabilities. Leaders in government and the private sector who share the responsibility for protecting the nation and its vital assets must face the stark reality that the risk of a terrorist attack is real and present almost everywhere.

The United States is an open society that offers terrorists many soft targets. Managing this risk requires the cooperative involvement of government, public, academia, and the private sector. Unfortunately, the risk of terrorist attacks cannot be completely eliminated. Therefore, the question is, what can we do to control and reduce this risk? The focus of this report is on the development and use of a methodology for making the right decisions to combat *catastrophic* terrorist attacks; specifically, attacks with potentially high consequences in terms of human, material, or financial loss. The study group concludes that the methods of quantitative risk assessment (QRA) based on the "set of triplets" definition of risk, best meet this challenge.

Qualitative methods, although certainly useful for screening potential terrorist attacks, do not provide the level of detail necessary for deciding on specific actions that provide the highest payoff in terms of preventing terrorist attacks, reducing their likelihoods, or reducing their consequences. On the other hand, quantitative methods have been explicitly developed for the purpose of assessing the risks of rare events that may have high consequences. They therefore readily apply to terrorism events.

Most QRA's adopt the "probability of frequency" framework for quantifying risk. In this framework, or model, the risk scenario of interest is assumed to occur over a long span of time with a certain frequency. Since we do not know the value of the frequency exactly, we express our knowledge in the form of a probability curve. This probability curve, obtained through the use of Bayes Theorem expresses our state of knowledge about the frequency, based upon all the relevant information, experience, and evidence available. Similarly, we use Bayes Theorem to probabilistically quantify the degree of damage resulting from the scenario. Thus, the same QRA approach, including the set of triplets, the probability of frequency, and the Bayesian treatment of the evidence applies equally well for terrorism scenarios as it does for "ordinary" risk scenarios.

Quantitative risk assessment methods are used in both government and the private sector. Current risk assessment practices are adequate for screening many attack scenarios, but the analysis demands are greater for assessing the risk from attacks that can have catastrophic consequences and have great uncertainties. Thus, the study group hopes that this report stimulates a serious dialogue about risk management applied to terrorist threats that could have catastrophic consequences.

The study group describes a scenario-based approach to QRA that has been used successfully to assess a wide range of risks, including those that fall within the purview of

government agencies such as the U.S. Nuclear Regulatory Commission, U.S. Environmental Protection Agency, U.S. Food and Drug Administration, and the U.S. Department of Agriculture. These agencies have used QRA to expose important issues that affect "making the right decisions," such as organizational and information structures. Information on terrorism may be limited, but if it is properly used and processed it can result in meaningful assessments of the risk of specific terrorist attacks. The information challenge is related more to gaining access to data than it is to the availability of data. Data are housed in thousands of databanks, both inside and outside the government, and serve a variety of end-users. Government at all levels, the private sector, academia, and the public are all important to national security and all have information sources important to combating terrorism. Therefore, they need to work in partnership. The private sector can contribute to upgraded security measures, work to improve analytical approaches, provide important information affecting possible terrorist attacks, and help identify serious infrastructure and technological vulnerabilities. The public must be made aware of the potential consequences of terrorist attacks and of the steps that can be taken individually and collectively to lessen those dangers.

Amorphous terrorist networks, such as al Qaeda, present a different problem than state-sponsored terrorism, because non-state terrorists thrive on mobility and have little fear of death or retaliation and locating them is extremely difficult. The result is the need for new ways of thinking about how to combat this new enemy. Advanced methods of risk assessment that have been developed explicitly to assess the risk of complex systems about which little is known hold the most promise for developing meaningful plans for combating terrorism.

Many national security decision makers consider such methods to be academic exercises with little relevance to the challenges they face; the study group concludes, to the contrary, that risk assessment and risk management methods can be particularly helpful for identifying the risks of, and vulnerabilities to, terrorist attacks, as well as for analyzing the likely effectiveness of proactive counterterrorist strategies. However, decision makers are most likely to use these tools if they are not overly complicated and if their utility can be demonstrated through scenario-driven assessments based on real-world examples of terrorist attacks.

In this report a general framework is presented of a methodology for risk assessment that could be adapted to analyzing the risk of specific terrorist attacks. The framework makes transparent the relationship between the components of the risk sciences, especially the link between risk assessment, decision-making, and management actions. The emphasis in the proposed methodology is "quantification", in the sense of making it clear that uncertainties have been factored into the results. The proposed methodology is supported by an experience base that provides confidence in its applicability to a wide range of risk issues, including terrorism. The methodology has a legacy of adaptability to many types of risk including health and safety, financial, and environmental impact. Also, the methodology is presented more in terms of general principles, basic definitions, and fundamental analytical concepts than as a "cookbook" for risk assessment.

The study group realizes that there are different ways to implement effective risk assessment and management. The approach presented here is an adaptation of an approach with a record of successful applications to a wide spectrum of risks ranging from mechanical failure

and human error, to sabotage, disease infestation, and catastrophic natural and manmade events. These applications included situations with limited availability of relevant data and loosely defined threats, resulting in the need to make *uncertainty* an inherent and explicit feature of calculating the risks. Finally, the methodology proposed is founded on several important tenets: a basic framework for combating terrorism, the "triplet" definition of risk, a specification of the requirements for analyzing terrorism risks, a set of guiding principles, and the steps necessary to carry out a thorough analysis.

THE BASIC FRAMEWORK

The basic framework for combating terrorism includes the following activities: (1) the collection and processing of intelligence on the intentions and capabilities of terrorists; (2) the processing of information on terrorist threats and target vulnerabilities; (3) identification, development, and analysis of the most likely terrorist attack scenarios; (4) decisions on actions to combat these scenarios; and (5) implementation of these actions to manage and minimize the risk of terrorist attacks.

The emphasis of this report is on activities 2 and 3. Based on available input from intelligence sources, risk experts carry out the type of analyses that could provide decision makers with actionable information on the assessed likelihood (probability) and linked uncertainties (confidence levels) of high-consequence terrorist attacks on particular targets.

THE TRIPLET DEFINITION OF RISK

The "triplet" definition of risk offers a basic structure for quantifying the risks. This structure has been used by government agencies and industries for more than three decades to analyze threats and vulnerabilities associated with a wide spectrum of man-made facilities and natural phenomena. The "triplet" definition indicates that, when we ask the question, what is the risk?, we are really asking three questions: (1) what can go wrong?; (2) what are the consequences?; (3) how likely are they? The first question is generally answered in the form of a structured set of scenarios that start with an initiating event (attack on system or target) and end with a consequence or a suite of consequences. The second question is answered by the various end states of the scenarios. The third question is answered by converting evidence about the scenarios into a measure of their likelihood (likelihood is interpreted in this report as a frequency with the uncertainty in the frequency being represented by a probability distribution).

REQUIREMENTS FOR ANALYZING TERRORIST ATTACKS

Terrorist attacks that could have catastrophic consequences impose special requirements on the method of analysis. The requirements are much more like requirements for analyzing threats exogenous to the target than for analyzing the types of risk that involve accidents and system failures. Exogenous threats include sabotage, earthquakes, volcanic eruptions, aircraft impacts, and other cataclysmic events, which we cannot predict with any certainty. Fortunately, because risk assessment has been used for exogenous events during the last few decades, we have an experience base to draw from in analyzing terrorist attacks. The experience base has

shown clearly that analyses of events with the characteristics of terrorism risk must as a minimum meet the following requirements.

- Quantification of uncertainty must be an inherent feature of the methodology.
- The methodology must be general, that is, applicable to all types of risk to maximize the lessons learned from the existing experience base in the risk sciences and to accommodate new methods as they evolve. Such generality includes the ability to calculate risk metrics that combine threats and vulnerabilities.
- The methodology must be supported by an experience base that demonstrates proof of concept.
- The methodology must be specialized for application to threats considered to have catastrophic consequences. Qualitative methods of risk assessment can be used to screen out attack scenarios where there is ample evidence that such screening is appropriate.

THE GUIDING PRINCIPLES

Guiding principles for scenario-based risk assessment have evolved through experience with actual applications. The principles presented below complement the above requirements.

- The quantitative expression of risk should be in the form of a structured set of scenarios, each having a corresponding likelihood and consequence.
- The set of scenarios must be complete in the sense that all of the important contributors to risk are included.
- The scenarios must be quantified in terms of clearly defined risk measures, must be realistic, and must incorporate uncertainties.
- Each scenario should depict a terrorist attack in the form of a sequence of events, starting with the initiating event that upsets an otherwise successful operation or system and proceeding through a series of subsequent events to the end state (i.e., the consequences of the attack). The initiating event must be based on a comprehensive threat assessment.
- Each scenario must accommodate combined events, including primary and diversionary events.
- The end-states must reflect initial, cascading, and collateral consequences (or levels of damage).
- Individual events and aggregated event uncertainties must be quantified on the basis of the evidence.
- The results must be ranked as to their contribution to risk in order of importance and must be presented in a way that supports decision-making.

These principles reflect important characteristics that have evolved from almost three decades of quantitative risk assessment work. Minor modifications have been made to specialize

these principles to the assessment of terrorism events. These principles are sufficiently general to provide the flexibility to accommodate different analytical tools. For example, different analysts may have different preferences for interpreting likelihood and probability.

TERMINOLOGY AND IMPLEMENTATION

Application of the methodology involves three distinct activities: threat assessment, system analysis, and vulnerability assessment. Threat is defined in this report as the intention of a terrorist to inflict harm or damage to a specific asset or target by a specific means or weapon system. System analysis is analysis of the target being attacked to determine how it functions, or how it can be degraded or destroyed. System has the usual meaning of an interdependent group of elements forming a whole to perform an intended function. A system may be a building, a complex of buildings and people, airline flight operations, an industrial plant, an entertainment complex, etc. Vulnerability is defined as the response of an asset or target to a terrorist attack, including the consequences of the attack. Thus, in this report the consideration of threats and vulnerabilities includes the consideration of targets, weapon systems, method of delivery, and consequences. Defining threat and vulnerability in this manner greatly facilitates the development of scenarios.

An initiating event for the scenarios (the attack on the system) is the output of the threat assessment. In this context, the initiating event for the 9/11 attacks was the terrorists taking over commercial airliners. All of the activities leading up to the taking over of the airliners are part of the threat assessment, the first of the three activities involved in the risk assessment methodology. A key aspect of the system analysis is establishing the baseline condition of the system or target to facilitate knowing when the system is in an upset or abnormal condition.

The vulnerability assessment is conditional on the results of the threat assessment and the system analysis, and involves the structuring of attack scenarios. The threat assessment provides the fundamental building blocks (the initiating events) for developing the attack scenarios. The system analysis defines the state of the system being attacked and is integral to the assessment of the vulnerability of the system to different levels of damage, including their consequences. The steps for implementing a total risk assessment of terrorist attacks are as follows:

1. Define the system being analyzed in terms of what constitutes normal operation and points of vulnerability to serve as a baseline reference point.
2. Identify and characterize the “sources of danger,” that is, the hazards (e.g., stored energy, toxic substances, hazardous materials, acts of nature, sabotage, terrorism, equipment failure, combinations of each, etc.).
3. Develop terrorist attack scenarios to establish levels of damage and consequences.
4. Adopt risk metrics that reflect the likelihoods of different attack scenarios in terms of target and collateral damage and quantify the scenarios based on the totality of relevant evidence.
5. Assemble the scenarios according to damage levels, and cast the results into the appropriate risk curves and risk priorities.

6. Interpret the results to guide the risk-management process.

HOW ALL THE PIECES FIT

This report is about analyses for making good decisions to combat terrorism. The focus is on assessing the risk of terrorist attacks in a manner that supports effective decision-making and terrorism risk management. The central theme of the report is an *analysis framework for terrorism risk assessment specialized to address high-consequence or catastrophic attacks for which there are limited data on the threat.* In addition to the methodology itself, the report addresses, in a limited way, the characteristics of the current terrorist threat and information and institutional issues affecting terrorism risk management. Assessment of the risk of terrorist attacks is clearly a "work in progress", an approach that is only one possible method, but with a history of diverse, successful applications in other realms and therefore provides a basis for further development.

The approach presented in the report is based on the principles and practices of QRA and has three basic parts: threat assessment, system analysis, and vulnerability assessment. The products of the threat assessment are the initiating events for the terrorist attack scenarios. System analysis defines the success and damage states of the system being attacked. Vulnerability assessment uses the output of both to structure scenarios connecting the initiating events to the damage states. Damage states include damage directly to the target, including its occupants, and secondary damage to the surrounding property and population and also possible extension or ripple effects with respect to economic and lifestyle impacts. The bottom line risk measure is the likelihood and uncertainty of different scenarios and their attendant consequences.

The *threat assessment* component of a terrorist attack QRA is a special challenge. It is an example of how the scientific process works—reducing observations to something that can be measured. Threats that have the potential of catastrophic consequences are hypothesized and deconstructed into the plans, resources, capabilities, and eventually the intentions of the terrorists. This kind of deductive logic modeling is common practice in the QRA field for quantifying initiating events. It requires the integration of intelligence information, experiential data, and an interdisciplinary group of experts. The presentation of results clearly communicates the uncertainties involved. This is achieved by defining the likelihood function in such a way that the uncertainties are clearly presented in the results and the results are explicitly anchored to the supporting evidence. Mapping a risk measure to the supporting evidence provides *objectivity* to the analysis.

For the other two components of the approach, *system analysis* and *vulnerability assessment*, a much greater experience base is available to draw upon. Both have been fundamental to the engineering analysis of structures and systems, including large dams, chemical complexes, transportation systems, petroleum production operations, and nuclear power plants. Thus, combined with threat assessment, the proposed methodology can meet the requirements set forth in the definitions, practices, and principles noted earlier.

The report includes an example of a limited-scope terrorist risk assessment. The example involves a simultaneous physical and cyberattack on a hypothetical electrical grid. The example

follows the above implementation steps and gives results in terms of selected damage levels to the grid and corrective actions for making the grid less vulnerable. As the example demonstrates, QRAs need not be long-term, multi-million dollar projects to provide important insights on the risks associated with complex electrical grids. The example assessment was performed in a matter of a couple of weeks with minimum resources.

OVERARCHING RECOMMENDATIONS

The report contains many recommendations for actions by federal and local government departments and agencies and by the private sector. The following five recommendations are of special importance. Recently (2003), the country's transmission grid came under increasing scrutiny because of events demonstrating its vulnerability. Appendix B provides an example of recent congressional testimony on this issue.

General Recommendation 1. Centers of excellence should be established for the study of quantitative risk assessment applied to the threat of catastrophic terrorism. The centers would provide a platform for research, development, and understanding of the terrorism threat and cutting edge mitigation strategies.

Currently much of the knowledge in QRA resides in the consulting sector, selected government agencies, and academia. Publications and journals are limited in this field. To engage the public in supporting these efforts financially and politically today, a more proactive public approach must be taken.

General Recommendation 2. The U.S. Department of Homeland Security should adopt policy guidelines for implementing a quantitative risk assessment process based on scientific principles that integrates the assessment of threats and vulnerabilities, clearly links the decision options with supporting evidence, and displays the characteristics of risks, benefits, and costs. The federal government should increase its investment in improving and refining these analytical tools for application to terrorism risk.

Modern analytical techniques developed for applications in other fields, such as nuclear power, environmental protection, and the chemical process industries, can be adapted for analyzing the threats posed by terrorists and the vulnerabilities of potential targets. These techniques can be particularly helpful for: (1) establishing priorities for the allocation of resources among many competing demands, and (2) ensuring that secondary effects on people and systems beyond the initial targets are adequately considered. The techniques can also provide insights into the effectiveness of alternative strategies for disrupting preparations for attacks by terrorist groups at home and abroad. The level of sophistication of the analysis should be commensurate with the risks involved.

General Recommendation 3. A priority of the U.S. Department of Homeland Security should be stimulating interest in sector-by-sector, structured risk assessments and related risk-reduction activities by federal and local governments, private-sector owners of targets, law enforcement organizations, and first responders. The objective must be to

ensure that mechanisms to address threats and vulnerabilities to terrorism are in place throughout the country.

An important lesson learned in the application of risk assessments is that risk is very site and situation dependent. The most important input to a risk assessment is often provided by the owner and operator of the asset that may be a target of opportunity. It is critical that the people closest to the target become engaged in the process of assessing threats and vulnerabilities. Their engagement not only provides assurance that the assessment is realistic, but also facilitates the ability to take preventive actions to decrease target vulnerabilities and to respond to an attack should one occur.

General Recommendation 4. The U.S. Department of Homeland Security should place the highest priority on the effective collection, fusion, and sharing of relevant data. The involvement of private-sector organizations will be essential, and consortia and other collaborative mechanisms, such as information sharing and analysis centers (ISACs), should be used whenever possible.

Reliable, timely information is the key to good decisions to counter terrorism at all levels of government, the private sector, and throughout the general population. Dozens of federal departments and agencies and hundreds of local agencies have extensive databases that contain important data, and the new department should be instructed and authorized to take steps to improve the mining of these databases for addressing immediate concerns and for developing long-term plans for protecting the nation.

General Recommendation 5. The federal government should provide incentives for the private sector to increase its investments in countering terrorist threats.

Some of the mechanisms that should be seriously considered are tax deductions for selected types of investments in private-sector security, adjustments in the antitrust laws that inhibit cross-institutional cooperation when such cooperation has broad national implications, and recognition through timely awards of pathfinding activities by private companies that increase the resilience of facilities and people to terrorist attacks. Attention should also be given to: (1) counterterrorism programs sponsored by professional societies and trade associations, and (2) initiatives at the local level to strengthen linkages between private-sector facility managers and law enforcement organizations. The federal government and the private sector should strongly support these collective efforts.

CHAPTER 1

A NATION CHALLENGED

On September 11, 2001, the United States was attacked, not by a rival state, but by a terrorist network. This single attack not only inflicted thousands of casualties and significant economic damage, but also profoundly changed the way Americans see themselves, their government, and their national security. The traditional dichotomy between domestic and international terrorism is no longer relevant. With improvised weapons of mass destruction and traditional weapons used in nontraditional ways, the power to cause harm has devolved from nation states to amorphous groups and even to individuals. Investments in national security at home and abroad must reflect these changing realities.

Because of the complexity of the issues and the rapid pace of change, the United States must use the best available analytical tools to identify, assess, weigh, and establish priorities for threats and vulnerabilities and to identify and evaluate options for action. The result is an urgent need for (1) understanding the threats involved, (2) appreciating vulnerabilities, and (3) an analytical process for assessing the risk and mitigating the threat. The risk sciences, developed for the purpose of assessing the risk of high-consequence, low-probability events in the presence of uncertainty, are considered the best approaches for developing a basis for decision making.

UNDERSTANDING THE THREAT

"Threat" has been defined as the potential intent to cause harm or damage to a system by adversely changing its state. However, in more general terms, it connotes an initiating event that can cause harm to a system or induce it to fail (Haimes and Horowitz, 2003). For example, improvised biological, chemical, and radiological devices that exploit technologies once the sole preserve of world and regional powers represent a significant threat. The proliferation of knowledge about improvised weapons of mass destruction has changed the nature of terrorism and elevated it to a strategic threat. In this conflict, the enemy wears no distinctive uniform, and there are no front lines and the U.S. "war" on terrorism is not a war in the conventional sense. This report suggests how the risk sciences can provide insights for the allocation of resources to combat this new type of threat.

UNDERSTANDING THE VULNERABILITY

Global strife and the perception of injustice create favorable conditions for terrorist organizations to recruit and raise support, but the execution of an attack depends upon the attractiveness of the target and the terrorists' resources and plan. To a terrorist, civilian populations; targets of historical, cultural, and national significance; and infrastructure that underpins the U.S. way of life are all "fair game." In this regard, the attackers on 9/11 exploited a systemic failure in the aviation transportation infrastructure to strike against both economic and military infrastructures.

As countries modernize, they become increasingly dependent on sophisticated technologies, with computers often controlling or linking vital, once disparate systems into national infrastructures that present unique targets to technologically sophisticated adversaries. Complex national infrastructures have critical nodes or choke points that, if attacked, could lead to significant disruption or destruction. Conventional assaults with truck bombs, dynamite, or cable cutting, as well as computer-generated attacks, could unleash a chain of events in which a service grid, an oil or gas pipeline, or an air traffic control system collapses with cascading effect.

Infrastructure vulnerability has to be thought of not only in terms of independent sectors, but also in terms of interdependent systems. A failure of the electrical power grid may affect not only the energy sector, but also in a cascading effect may result in the collapse or severe disruption of transportation, telecommunications, public health, and banking and financial systems. In Chapter 4 of this report, a limited scope risk assessment is performed on a hypothetical electrical grid to illustrate that such risks can be analyzed. The example shows how the results can then be used to guide actions for reducing risk and cascading impacts.

RESPONDING TO THE CHALLENGE

To meet the challenge of terrorism, new analytical tools and new institutional arrangements must be developed. *Making the Nation Safer*, a report by the National Research Council (NRC, 2002), recommends that the United States use existing technologies and initiate research and development in a number of critical areas to aid the nation in its war on terrorism. The question asked in the present study is what technologies are available right now. The study group believes that the appropriate application of the risk sciences is a definitive way to begin improving homeland security.

The recently created U.S. Department of Homeland Security has developed *The National Strategy for the Physical Protection of Critical Infrastructures and Key Assets* (DHS, 2003). One of the eight guiding principles that underpin this strategy is "Develop technologies and expertise to combat terrorist threats." Another is that the federal government must provide and coordinate "national-level threat information, assessments, and warnings that are timely, actionable, and relevant to state, local, and private sector partners." The S*trategy* contains numerous initiatives, two of which are (1) to identify key protection priorities and develop appropriate supporting mechanisms for these priorities, and (2) to encourage sharing of risk management expertise between the public and private sectors.

Even before 9/11, there were many calls for using the risk sciences to combat terrorism. In April 1998, the General Accounting Office cited the growing use of risk assessment in both the public and private sectors to support decisions for prioritizing security investments (GAO, 1998). Other agencies, such as the U.S. Coast Guard, had taken steps before 9/11 to make greater use of the risk sciences to manage the risk of both marine accidents and terrorist attacks (Garrick, 1999). More recently the Coast Guard published their *Risk-Based Decision Making Guidelines* (USCG, 2001). The National Aeronautics and Space Administration (NASA) published a procedural guide on probabilistic risk assessment for NASA managers and practitioners (NASA, 2002). The U.S. Department of Transportation Research and Special

Programs Administration (RSPA) utilizes risk management concepts and tools to prioritize compliance activities and address the risks associated with noncompliance. RSPA has released a report of a risk management framework for the transportation of hazardous materials (ICF Consulting, 2000).

Other institutions that advocate using risk-management technologies to enhance decision-making are the National Academies, various academic institutions, and the private sector bodies, including not-for-profit think tanks. The National Academies have long tracked the evolution of the risk sciences and their contributions to society. The National Academies has published a series of studies starting with a report in 1983 on risk assessment in the federal government (NRC, 1983). This was followed by a report on improving risk communication (NRC, 1989), a report on risk-informed decisions in a democratic society (NRC, 1996a), and several reports recommending increased use of risk assessment and the risk sciences in the environmental remediation of the nation's nuclear laboratories (NRC, 1994a), chemical weapons disposal (NRC, 1994b), the management of transuranic radioactive waste (NRC, 1996b), and the safety assessment of the space shuttle (NRC, 1988).

Several academic centers have advanced the use of the risk sciences in health, natural hazards, engineered systems, finance, and economics. Among these are the University of Virginia Center for Risk Management of Engineering, the Harvard Center for Risk Analysis, and Clark University George Perkins Marsh Institute. Many university engineering schools have very strong programs in the risk sciences, including the Stanford Management Science and Engineering Department, the Carnegie Mellon University Engineering and Public Policy Department, and the Massachusetts Institute of Technology Nuclear Engineering Department. Numerous business schools have active programs involving the risk sciences, one example of which is the Wharton School of the University of Pennsylvania.

It became clear from the briefings to the study group that although risk assessment is an established discipline throughout government, academia, and the private sector, the actual amount of experience of using formal, comprehensive, and quantitative methods of risk assessment to analyze the risk of terrorism is fairly limited. The study group drew a sharp distinction between *qualitative* and *quantitative* methods of risk assessment. It is clear that analyzing the risk of terrorist attacks with the potential for catastrophic consequences requires methods of analysis that systematically and rigorously quantify uncertainties. Qualitative methods can be used to screen the risks of terrorist attacks, but much more is required to quantify the risk of genuine threats that have potentially catastrophic consequences.

The method of quantitative risk assessment presented in this study is one way of doing such analyses. While the experience base for the quantitative methods proposed in this study is extensive, it is so in very select areas. Clearly, the nuclear power industry has championed the development and application of quantitative methods more than any other industry or sector. Taking the risk sciences as a whole, the use of quantitative methods is quite limited. In fact, it suggests the need for the creation of "think tanks" dedicated to quantitative methods and applications of risk assessment. In particular, centers of excellence should be created where the study of QRA applied to terrorist-inspired catastrophes could become the subject of study and analysis with application in both public and private sectors.

QRA AND THE OVERALL ANALYTICAL FRAMEWORK FOR ACTION

The context for QRA in the overall framework for action to counter a terrorist threat can be presented as follows:

Step 1. Intelligence gathering (intentions and capabilities of terrorists).

Step 2. Information processing (threats and vulnerabilities).

Step 3. Identification, analysis, and development of the most likely terrorist attack scenarios, including their consequences (based on the evidence from Steps 1 and 2).

Step 4. Decision-making on actions to combat terrorist attacks.

Step 5. Implementation of actions.

This report considers all five steps, but emphasizes Step 3 based on the belief that high-quality analyses of the risks of terrorist attacks are the most important input to making the right decisions. Steps 1 and 2 are addressed primarily in the context of understanding the threat and information requirements to support meaningful analysis of the terrorism risks. The study group believes Step 3, QRA, is the best course to take to support effective decision making on the most serious threats. Step 4 is addressed by illustrating the connection between risk assessment and decision analysis. Finally, Step 5 is addressed in the context of organizational issues associated with implementing actions to combat terrorism.

Intelligence Gathering and Information Processing (Steps 1 and 2)

Steps 1 and 2, observations, information, and data from a variety of sources are converted into a quantitative, numerical form suitable for modeling. The basis for understanding risks is supporting evidence. To quantify the risk of terrorist attacks, it is necessary to have not only intelligence and other information, but also to structure information in a form suitable as input to a QRA and to the subsequent decision analyses (Step 4). The questions addressed in the intelligence gathering and information-processing steps are: which threats are considered the most serious, and what is the supporting evidence for those threats? Answers to these questions must be in the form of targets, weapons, and delivery systems. This evidence is the basis for a first-pass screening and prioritizing of attack scenarios and an identification of the scenarios on which we need to turn up the microscope in Step 3. Intelligence and information experts, as well as experts on the weapons and delivery systems, can cast this information into a form suitable for QRA and subsequent decision analysis.

Identification, Analysis, and Development of the Most Likely Attack Scenarios (Step 3)

The first two steps provide a basis for identifying threats that should be analyzed in detail. They screen out the less important threats so resources can be concentrated on the more serious, more credible threats. The methodology presented in Chapters 2 and 3 is a structured example of this presented as a three-part process: threat assessment, system analysis, and vulnerability assessment.

We note that in Step 3, threat assessment includes information not only on the intentions and capabilities of the terrorists, but also information on targets and weapon delivery systems. System analysis refers to the system being attacked and the need to define successful operation of the system as a baseline for knowing how the system can fail or be destroyed. Vulnerability assessment is the response of the system to the threat and includes consequences.

Decision Making and Implementation of Actions (Steps 4 and 5)

Decision analysis involves determining the risks, costs, and benefits of different alternatives available to a decision maker. While a risk assessment is not a decision analysis, the linkage between the two for high-consequence events is such that good decisions are very strongly dependent on an understanding of the risks. The making of decisions is followed by actions. It is much easier for decision makers to have the support of the public if those actions are supported by strong evidence—evidence that has been put through a transparent systematic risk and decision analysis process as outlined in Step 3 above.

Chapters 2 and 3 present an introduction to a QRA methodology for determining risks from different terrorist attack scenarios. An important question is: when do we do a QRA of the type presented in Chapters 2 and 3? The answer is: when it is important to making a decision about actions to counter terrorism. In reality, not all actions need to be based on a QRA. Realistically, some screening and sorting out of the risks can be done at the intelligence and information processing levels. As indicated earlier, QRAs are generally reserved for the threats with potentially catastrophic consequences when preliminary analysis indicates a reasonable likelihood of occurrence. The likelihood may be low, but the consequences of the event may be so catastrophic that it needs special attention.

CONCLUSION AND RECOMMENDATION

Conclusion. Risk assessment is an established discipline of the management sciences, although applications of quantitative and rigorous models are limited.

Recommendation. Centers of excellence should be established for the study of quantitative risk assessment applied to the threat of catastrophic terrorism. Their purpose would be to provide a platform for research, development, and understanding of threats and to provide cutting-edge mitigation strategies.

REFERENCES

DHS (U.S. Department of Homeland Security). 2003. The National Strategy for the Physical Protection of Critical Infrastructures and Key Assets. Washington, D.C.: U.S. Department of Homeland Security.

GAO (Government Accounting Office). 1998. Combating Terrorism: Threat and Risk Assessments Can Help Prioritize and Target Program Investments. GAO/NSIAD-98-74. Washington, D.C. General Accounting Office.

Garrick, B.J. 1999. Risk assessment methodologies applicable to marine systems. Marine Safety Council Proceedings: The Coast Guard Journal of Safety at Sea 56(3): 50-52.

Haimes, Y.Y., and B. Horowitz, 2003. Adaptive Two-Player Hierarchical Holographic Modeling Game for Counterterrorism Intelligence Analysis. Submitted to Systems Engineering.

ICF Consulting. 2000. Risk Management Framework for Hazardous Materials Transportation. Prepared for the U.S. Department of Transportation. DTRS56-99-D-70123. Fairfax, Va.: IFC Consulting.

NASA (National Aeronautics and Space Administration). 2002. Probabilistic Risk Assessment Procedures Guide for NASA Managers and Practitioners. Prepared for Office and Mission Assurance. Washington, D.C.: National Aeronautics and Space Administration.

NRC (National Research Council). 1983. Risk Assessment in the Federal Government: Managing the Process. Washington, D.C.: National Academy Press.

NRC. 1988. Post-Challenger Evaluation Space Shuttle Risk Assessment and Management. Washington, D.C.: National Academy Press.

NRC. 1989. Improving Risk Communication. Washington, D.C.: National Academy Press.

NRC. 1994a. Building Consensus through Risk Assessment and Management of the Department of Energy's Environmental Remediation Program. Washington, D.C.: National Academy Press.

NRC. 1994b. Recommendations for the Disposal of Chemical Agents and Munitions. Washington, D.C.: National Academy Press.

NRC. 1996a. Understanding Risk: Informing Decisions in a Democratic Society. Washington, D.C.: National Academy Press.

NRC. 1996b. The Waste Isolation Pilot Plant: A Potential Solution for the Disposal of Transuranic Waste. Washington, D.C.: National Academy Press.

NRC. 2002. Making the Nation Safer: The Role of Science and Technology in Countering Terrorism. Washington, D.C.: The National Academies Press.

USCG (U.S. Coast Guard). 2001. Risk-Based Decision Making Guidelines. 2nd Ed. Available online at: http://www.uscg.mil/hq/gm/risk/intro.html

CHAPTER 2

OVERVIEW OF QUANTITATIVE RISK ASSESSMENT (QRA)

RISK MANAGEMENT IN BRIEF

Risk-management analysis technologies include methods of quantitative analysis that can be applied in the assessment of terrorist-initiated events. The study group believes that the risk-management sciences can provide a basis for allocating investments to improve homeland security. Risk-management techniques have been successfully applied in other fields to assess technological risks and external threats and low-probability, high-consequence events. In this chapter, the basic concepts of risk management for assessing terrorist threats and infrastructure vulnerabilities are introduced. Chapter 3 addresses some of the more analytical concepts involved in the quantification process of risk assessment.

Risk management is based on established principles, and includes the following activities (Haimes, 1998):

1. Risk assessment, an objective and preferably quantitative evaluation of risks, including threats and vulnerabilities.
2. Risk communication, the dissemination of risk information to stakeholders in an understandable form.
3. Decision analysis, a determination of appropriate "corrective actions" or mitigating measures to reduce the risks.
4. Risk mitigation, the implementation of corrective actions based on the decisions.

The study group believes that prudent application of risk-management practices to possible terrorist attacks can help in many ways: (1) by making targets less attractive to terrorists, thereby reducing the likelihood that those specific targets will be attacked; (2) by lessening damage to the target in the event it is attacked; (3) by reducing the time necessary for recovery from an attack; and (4) by lessening the effects of collateral and cascading damage in areas surrounding a target. Combating terrorism effectively requires an understanding of threats, reliable information, teamwork on the part of many segments of society, organizational entities capable of implementing actions, and good supporting analyses. An example of a successful risk-management program is provided in Box 2-1.

Box 2-1. Risk-Management Aspects of Fire Safety

Fire safety is an example of risk management in action. Consider a fire chief in a small jurisdiction. With limited funds and his own creativity, the chief must decide how he can reduce both the number of fires and the impact of fires that do occur. There are a number of potential scenarios (e.g., different origins of fires, such as domestic carelessness, industrial accidents, natural disasters, and arson), and numerous corrective measures could be implemented (e.g., fire safety promotion, more fire crews, and more training for security personnel).

In evaluating the scenarios, the chief is conducting a *risk assessment*. Considerable *evidence* can assist him in this process. Past accidents, catastrophes, near misses, and instances of sabotage can help him assess both probabilities and end-states for different scenarios.

Risk communication is the interaction between those carrying out the assessment (e.g., the chief) and those who make decisions and implement changes (e.g., the public). If, for instance, the risk assessment indicates that smoking in bed is a serious problem, passing this information on in a form that facilitates decision-making is essential. The risks to people who choose to smoke in bed should be highlighted. What causes their habit of smoking in bed? How can they be interested in fire safety? What mitigating measures can be taken in the event of fire?

Once scenarios have been mapped out, a *decision analysis* can be done. The fire chief and others determine implementable actions that will have the most impact on reducing the number and consequence of fires.

Finally, *risk mitigation* measures are taken. Appropriate preventive measures are implemented, monitored for their success in reducing risk, and revised as necessary.

This example illustrates the localized nature of assessing vulnerabilities and implementing risk mitigation measures. Fire prevention has been achieved at the national level through improved building codes that mandate fire escapes, escape routes, sprinklers, barriers, and other measures. This underscores the need for federal, state, and local cooperation.

Risk management for fire safety is promoted by the National Institute of Standards and Technology and other institutions, based in large measure on lessons learned from actual fires. Activities include evacuation and other practice drills, investments in people and equipment (by both owners and local entities), addressing shortcomings (e.g., the proximity of hydrants and firehouses), identifying triggering events, and inspecting electrical systems and the storage of flammable materials. Careful risk management has not only saved lives, but has also yielded financial benefits for property owners, the insurance industry, and the larger community of ordinary citizens.

A Definition of Risk

All aspects of risk management are essential to controlling risk, but the foundation for making the decisions is knowing what the risks are. The science of risk assessment, particularly quantitative risk assessment, is an analytical process designed to answer three basic questions about risks from a system point of view (Kaplan and Garrick, 1981):

1. **What can go wrong?**
2. **How likely is that to happen?**
3. **What are the consequences if it does happen?**

These questions, known in the risk sciences as the "triplet definition of risk," provide a general framework for all types of risk assessment. The triplet definition of risk is covered in a later section, and a brief history of risk assessment is presented in Appendix A.

Many government agencies and organizations in the private sector use the triplet definition. Examples include the U.S. Nuclear Regulatory Commission (USNRC, 1999), the National Aeronautics and Space Administration (NASA, 2002), the U.S. Department of Energy (DOE, 2001), and most of the nuclear electric utilities. The National Research Council has referenced or recommended the "set of triplets" definition in numerous studies on risk, including reports on the Challenger space shuttle accident (NRC, 1988), the disposal of chemical weapons (NRC, 1994b), the environmental remediation of the nation's nuclear laboratories (NRC, 1994a; NRC, 1999), the management and disposal of nuclear waste (NRC, 1996b), and a topical report

on understanding risk (NRC, 1996a). Among other agencies referencing or applying the definition are the U.S. Department of Energy (Helton et al., 2000), the U.S. Department of Defense (Johnson et al., 1997), the North Atlantic Treaty Organization (Garrick and Kaplan, 1995), and the U.S. Coast Guard (Garrick, 1999).

Risk Communication

Risk communication has been defined by the National Academies as "an interactive process of exchange of information and opinion among individuals, groups, and institutions. It involves multiple messages about the nature of risk and other messages, not strictly about risk, that express concerns, opinions, or reactions to risk messages or to legal and institutional arrangements for risk management" (NRC, 1989). The purpose of risk communication is to raise the level of understanding of issues and actions having to do with risk. The key words are a "process of exchange of information," "opinion," and "understanding." Risk communication is an essential part of decision making. Exchanging information to facilitate an understanding of risks enhances the process of developing consensus on issues and taking action in the context of risk management "best practices."

Decision Analysis

Decision analysis is a well-established discipline that is based on the principles and practices of decision theory (Raiffa, 1996). In fact, there is a long history of university business schools using risk and decision theory tools, such as Bayesian analysis, to develop forecasting models for finance and economics. The risk sciences have drawn on this experience to develop methods for using limited information to quantify the likelihood of catastrophic events. Decision analysis as a subset of decision theory has been formally recognized since 1963 (Howard and Matheson, 1989). Decision analysis considers the three major attributes associated with all decisions: risks, benefits, and costs.

Difficulties often encountered in decision analysis include: (1) major uncertainties about the outcomes of decisions; (2) system dependencies; (3) time horizons that vary from nanoseconds to geological times; and (4) "value judgments," such as on the importance of civil liberties, and the need to balance the various forms and combinations of costs, risks, and benefits. The demand for openness in public decision-making has encouraged the adoption of formal, explicit, visible, systematic and quantitative methods of dealing with risk.

This report is about using quantitative risk assessment to help make the "right decisions" on how to combat terrorism. Therefore, it is important to establish the link between QRA and decision analysis. To do this, consider Figure 2-1, a diagram of a typical decision problem. At the point of decision x we are faced with a choice between a set of options $\varphi_1 \ldots \varphi_x$. If we choose option 1, the resulting outcome is shown as R_1. If we choose option φ_2, the result is R_2 and so on. Now, it may be (and usually is) the case that at the point of decision, x, we are uncertain about the outcomes R_i. In that event we express our uncertainty by expressing the outcomes as probability curves as in Figure 2-2.

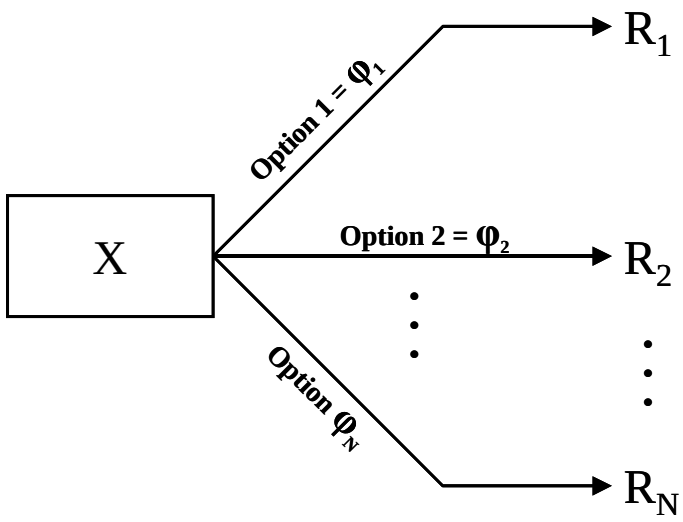

Figure 2-1. Fundamental Decision Diagram

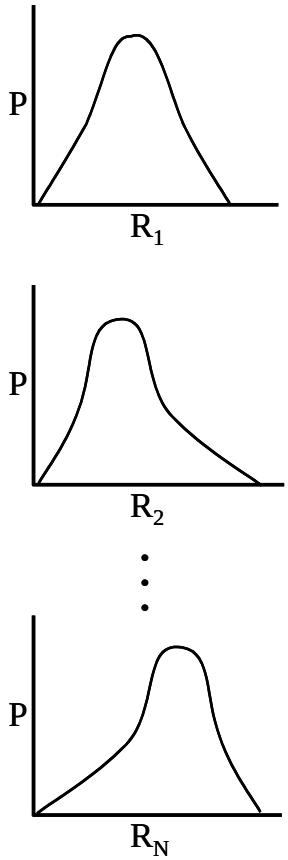

Figure 2-2. Expressing Uncertainty About the Outcomes R_i

Suppose now that the outcome of an option choice is not a single number, but instead is a vector quantity, reflecting the results of a decision in terms of money, health, fame, or

whatever. To include this case in our analysis, we simply regard the outcomes $R_1 \ldots R_n$ as outcome "vectors," e.g.,

$$R_n = \begin{bmatrix} R_{n,1} \\ R_{n,2} \\ \vdots \\ R_{n,k} \end{bmatrix}$$

In making the decision now, we simply choose that φ_n that gives us the most favorable outcome vector R_n.

If the vectors R_n are uncertain, we express that uncertainty with probability curves over the vector space. At the present time we omit the details of this process as unnecessary and distracting to the reader. However, we symbolically reflect the results of this process in Figure 2-3, which shows the outcome vectors expressed as probability curves reflecting our uncertainty.

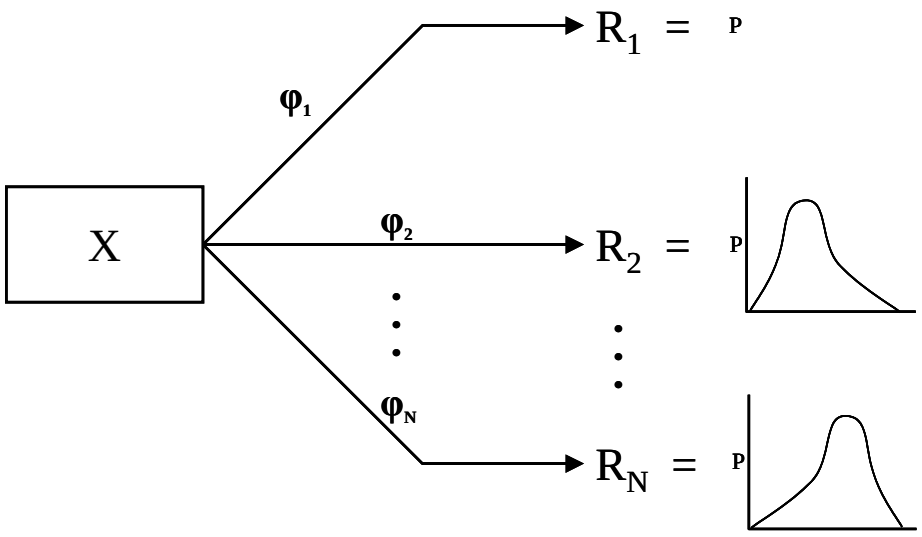

Figure 2-3. Decision Diagram When the Outcome Vectors are Uncertain

The purpose of QRA is to produce the probability curves in Figure 2-3. Given these curves, the decision is now made by choosing that option, call it φ^*, that gives us the most favorable outcome.

The curves represented in Figure 2-3, which are the outcomes of the QRA, should reflect the whole body of evidence available at the time the QRA is done. Should new evidence become available, the QRA should be updated, yielding new curves and possibly a new decision.

A GENERAL FRAMEWORK FOR THE QRA OF TERRORIST-INITIATED EVENTS

The triplet definition of risk, a set of guiding principles, and the risk assessment steps of the methodology presented in this report have been structured as a general approach to risk assessment applicable to terrorism risk as well as to other types of risk. The methodology: (1) provides a structure that can accommodate a variety of detailed methods that are available for risk assessment; and (2) provides flexibility for its application to any type of risk, including the risk of terrorism.

The methodology has not evolved in the abstract, but through applications in such areas as chemical, petroleum, nuclear, explosives, wastes, transportation, and space; and for risk challenges such as animal importation, security, food safety, dam safety, earthquakes, fires, water resources, and severe storms. Five textbooks are referenced that contain many examples of the diverse application of scenario-based risk assessments (Lewis, 1990; Haimes, 1998; Molak, 1997; Fullwood, 2000; Blockley, 1992). Comparisons have been made of how risk assessment is performed in different industries using the principles outlined in this report (Garrick, 1988, 1989a; NRC, 1988).

An assessment of security risk to the space shuttle, which included a consideration of terrorist attacks, was performed prior to 9/11 using the same methods as are presented here (Garrick and Kaplan, 1999). The National Aeronautics and Space Administration has several ongoing risk assessment activities using the same principles (Fragola, 1995; Woods, 1997). Efforts are also ongoing to apply the methodology to making decisions about disarmament (Garrick and Kaplan, 1995). The U.S. Coast Guard is actively adopting risk assessment technologies (Garrick, 1999). The nuclear industry has the most experience—with risk assessments for most of the more than 400 nuclear power plants throughout the world; and the risk assessment experience of U.S. nuclear plants using the same methodology is well documented (Garrick and Christie, 2001; Garrick, 1989b).

A key point is to demonstrate how threats and vulnerabilities are integrated. Elements of the integration process are shown in the sample application in Chapter 4. However, this example does not model all of the initial stages of the threat assessment. A threat is defined as "an indication of something impending, or an expression of intention to inflict evil, injury, or damage." This definition is extended in this report to *the intention of a terrorist to inflict harm or damage to a specific asset or target by a specific means or weapon.* The reason for defining threat this way is to facilitate the development of attack scenarios, which are bounded in terms of the intentions and capabilities of the terrorist. For similar reasons, we have chosen to define vulnerability as *the response of an asset or target to a terrorist attack, including the consequences of the attack.* Both definitions include not only the intentions of the terrorist and the response of the specific targets being attacked, but also the weapon, the delivery system, and the consequences.

While there is much more experience with quantifying vulnerability risk than the quantitative risk assessment of threats, the methodology presented in this report has been structured to include both. The methodology involves three major steps: (1) analyzing the threat, (2) characterizing the success state of the system under attack—or the success scenario in the methodology, and (3) the vulnerability assessment. The threat analysis generates the initiating

events for the vulnerability assessment. As with all risk assessments, the initiating events are application dependent and require extensive involvement of experts—those who develop and analyze intelligence and those who are expert in the nature of the threat, for example bioterrorism.

The approach includes a threat model having the form of an incoming logic tree, whose undesirable event or "top event" becomes the initiating event for assessing the vulnerability of the system being attacked. The approach has been widely used in assessing the risk of nuclear power plants (Garrick, 1989b). In other words, *the output of the threat assessment is the input to the vulnerability assessment.* Suppose we have intelligence that rockets might be used to deliver a large quantity of nerve agent to a major outdoor spectator event such as a football game of high national interest. How do we go about using the proposed methodology to quantify the risk of such an event? The first thing we do is "define the system." In this case we examine the stadium or stadiums to establish how the stadium operates, the most likely points of attack, possible staging locations for added protection, what safety equipment exists, evacuation routes, etc. This becomes the basis of the success scenario. Second, we "characterize" the hazard, in this case the threat of a nerve gas attack, on the basis of evidence provided by the terrorist experts. Third, we develop terrorist attack scenarios with possible extensions. We need to establish the initiating events for the different scenarios, that is, the scenarios that result from attacking the football stadium—or the different pathways through our event trees (discussed in the next chapter).

The initiating events result from the threat assessment. The top event for the threat assessment in this case could be the delivery of "X" amount of nerve gas to the stadium (based on intelligence sources that have determined this is the threat of greatest immediate concern). The likelihood, method of delivery, and general target of the attack would come from the same sources and knowledge experts. The exact point of the attack would come from assessments by military and facility experts of how the attack could cause the greatest amount of damage based on the method of delivery and type of weapon.

The goal then is to convert intelligence information into some sort of numerical form that can be used in a threat assessment model. How does one go about doing this? How does one take uncertain information that is not numerical and convert it into an input parameter for a logic model? The answer is, in the same way it has been done for many similar risk assessments in the past. First, a framework for processing the information is developed. The framework often used is to think in terms of events per unit of time, that is a frequency. The event in this case is a nerve gas attack on a stadium—an attack that has never happened. But the evidence indicates it may happen now. So, we can suppose that the evidence will exist for the next 100 or 1,000 years or so and ask, based on the current evidence and if nothing changes, how often would we expect such an attack to occur? Thus, we have introduced the notion of time into our thought experiment. Exercising the thought experiment with various experts (intelligence, weapons, delivery systems, information, etc.) and examining their underlying evidence can enable us to connect the initiating event to a knowledge-based frequency.

The experts and their knowledge base (i.e., the supporting evidence for their opinions), become the basis for assigning an objective probability distribution to the frequency of the

attack. The supporting evidence is critical to assigning these probabilities. The evidence is quantified by representing it as a probability distribution that clearly communicates the uncertainty based on the quality of the evidence. Low-quality evidence means that the input distributions and, therefore, the uncertainties have a wide spread to them. Including the uncertainties is an essential part of "quantification." Considering the uncertainties in each scenario and aggregating the scenarios leads to quantifying the total risk. Several different types of threats are discussed below.

Box 2.2 Quantification of Threats

The key to quantifying the threat (target, weapon, and delivery system) of a terrorist attack is being able to account for varying levels of uncertainty about the likelihood of the threat. Quantification of likelihood means that the threat is represented by a mathematical parameter that embodies no more and no less than what can be supported by the evidence. A mathematical parameter that has been successfully used to represent the likelihood of rare and uncertain events is "probability of frequency." This parameter has been effectively used in quantitative risk assessments having the same circumstances found in threat assessment—bits and pieces of evidence but no direct frequency data (several examples are in the references at the end of this chapter). Many methods exist for converting "bits and pieces" of evidence into parametric forms that are accountable to varying degrees of uncertainty. The methods, for the most part, are rooted in expert elicitation and inferential analysis (Bayes Theorem); both are established methods of converting information to a form suitable for quantitative analysis.

There are several keys to the process of converting limited information to quantifiable parameters. The first and foremost is to embrace uncertainty as a fundamental part of the analysis. This has been the primary driver for the methodology presented in this report. Others have also highlighted the importance of making uncertainty a part of the quantification process of analysis. Quoting from Heuer (1999) from the Center for the Study of Intelligence, "Managers of intelligence analysis need to convey to analysts that it is okay to be uncertain, as long as they clearly inform readers of the degree of uncertainty, sources of uncertainty, and what milestones to watch for that might clarify the situation. Inserting odds, ratios, or numerical probability ranges in parentheses to clarify key points of an analysis should be standard practice." This is exactly what "probability of frequency" risk parameters do in a formal way—display the uncertainties and link them to the supporting evidence.

A second key point is to involve multiple experts. It is not enough to leave it to the opinions of any one set of experts. There must be interaction and exchange with experts on the threats involved, experts on eliciting information, experts on inferential analysis, and experts on risk assessment. For example, if the threat is bioterrorism, then experts on biological agents, biological weapons, and delivery systems are essential.

A third key input is a disciplined process for identifying, structuring, and prioritizing threat scenarios. There are many methods for doing this, one of which is advocated by Heuer. Referred to as the analysis of competing hypotheses (ACF), it is based on "an eight-step procedure grounded in basic insights from cognitive psychology, decision analysis, and the scientific method." Replacing the word "hypothesis" with the word "scenario" in the ACF process generally replicates the approach that has been used extensively in environmental and nuclear safety applications under the general descriptor of "expert elicitation." Many formal methods of eliciting information have been successfully applied (Hora, 1991; Kaplan, 1992; LLNL, 1995; USNRC, 1996).

Nuclear and Radiological Attacks

The nuclear and radiological threats are nuclear weapons (existing or improvised), so-called "dirty bombs" using conventional explosives to spread radioactive material, and attacks on nuclear power plants and fuel storage facilities. Applying the risk assessment methodology of this report would involve obtaining intelligence information on each of the threats from the intelligence experts and transforming that information into (1) the most likely targets and (2) the basic events of the threat model as described above. Much of the nuclear industry is in better shape than most because of the extensive amount of risk assessment work that has already been performed. The scenarios are often well defined and the vulnerability of most nuclear facilities has been well analyzed. In the case of nuclear power plants and many other nuclear facilities, these scenarios already exist. The major requirement for analyzing the risk of such attacks would be in the development of initiating events based on a threat assessment.

Bioterrorism Attacks

A National Research Council report identified two types of biological terrorist threats: (1) communicable infectious agents, and (2) "biological agents that may cause disease or death in individuals but generally may not be transmitted between individuals" (NRC, 2002). Examples of the first type are smallpox, Ebola, and foot-and-mouth disease; an example of the second type is anthrax. How might the threat of a bioterrorism attack be addressed in the context of the proposed methodology?

It is possible to systematically assess the threat of a bioterrorism attack by using existing threat assessment methodologies. It is possible because, over the years, data have been collected from incidents involving such threats. Such data include (in addition to actual bioterrorism events and natural events, such as the outbreak of West Nile virus and the spread of anthrax spores in the United States and the foot-and-mouth disease in the United Kingdom) the analysis of threat statements from terrorist groups, profiles of perpetrators, interviews of acquaintances, and patterns of inquiries by candidate perpetrators on the Internet, for example. There also exist methodologies for accomplishing what might be called preliminary threat assessments, preliminary in the sense that such results provide important input to the threat model of the proposed methodology of this report. It has been pointed out, "that none of the past known biological or other terrorist attacks that caused mass casualties were preceded by the issuance of any warning." The study group believes that it is not so much a matter of "no warning" as it is a lack of access to data and information by persons with the necessary expertise to see the warnings and precursor events as discussed in Chapter 5.

Toxic Chemicals and Explosives

Classes of chemicals that have been identified as candidate terrorist weapons are (1) chemical weapons developed for military applications, (2) toxic industrial chemicals, and (3) explosives and highly combustible materials.

A prime candidate weapon for a terrorist would be chemical weapons because of their general availability and advanced development; many well-tested delivery systems are also available. Building risk models of chemical weapon attacks have been greatly facilitated by the risk assessments that have been performed supporting the nation's chemical stockpile disposal program. While these models were not developed for the purpose of analyzing the risk of terrorist attacks, they did address the threat of external events such as earthquakes and aircraft impacts. When combined with internal event risk assessments, these models provide an excellent base for models of terrorist threats.

The United States produces, transports, and stores large quantities of toxic industrial chemicals. Terrorist attacks on industrial plants, storage sites, and pipelines could release toxic chemicals such as volatile acids, chlorine, and phosgene in dense population centers. The releases could come about by a variety of means including deliberate actions by insiders, explosive charges, or severe damage to pipelines and storage tanks.

Transportation systems with chemical cargo could become targets in maximum impact locations. Risk assessment-type studies have been performed for many of those facilities and transportation systems in the U.S. that provide a reasonable evidence base of their vulnerability

to accidents and external threats, such as earthquakes, exterior fires, and severe storms. To extend these assessments to terrorist risk assessments would require intelligence on terrorist threats that might target such facilities. The chemical plant and transport experts would be the best sources of information on points of vulnerability. These sources of information, together with the risk and safety studies that are available on the targets of opportunity, would be excellent starting points for identifying the basic initiating events and their likelihoods for the threat assessment. Existing risk and safety studies would go a long way toward developing the risk assessment components of a terrorist attack, the scenario, and the vulnerability of specific targets.

Considerable evidence exists on the use of explosives and combustible agents to carry out a terrorist attack. Besides the attacks of 9/11 there have been some 10 major events, and more than 2,000 smaller events against the U.S. that have involved explosives and combustible agents. These attacks resulted in more than 1,000 fatalities and 6,000 injuries. The attack on the Murrah Federal Building in Oklahoma City was one such attack. This type of information and other intelligence information would be the starting point for analyzing the risk of attacks using explosives and flammable materials. As before, the idea would be to choose those targets of opportunity as determined by the available evidence, including intelligence information, and construct incoming threat master logic diagrams of the type described in the next chapter. The targets of greatest opportunity would be the basis for selecting the top event of the threat fault tree. The basic initiating events would be based on the combination of experience to date with such attacks and the intelligence information.

A Cyberattack and Physical Attack on Critical Infrastructure

The sample application in Chapter 4 demonstrates the principles and assessment steps of the methodology, but is not a comprehensive risk assessment. On the back end, the consequence analysis ended with calculating different damage states of the critical infrastructure rather than proceeding further with an analysis on economic, health, and safety aspects. It was limited on the front end by considering many initiating events for the attack, but did not consider the basic initiating events themselves. The decision to limit the scope of the analysis was not due to any limitation to the methodology, but by choice was on the basis that our goal was only to demonstrate how the methodology could be applied.

REFERENCES

Blockley, D., Editor. 1992. Engineering Safety. London, U.K.: McGraw-Hill International.

DOE (Department of Energy). 2001. Total System Performance Assessment Site Recommendation: Yucca Mountain Project. Washington, D.C.: U.S. Department of Energy.

Fragola, J.R. 1995. Probabilistic Risk Assessment of the Space Shuttle: A Study of the Potential of Losing the Vehicle During Nominal Operation. SAICNY95-02-25. Prepared for the National Aeronautics and Space Administration.

Fullwood, R.R. 2000. Probabilistic Safety Assessment in the Chemical and Nuclear Industries. Boston, Mass.: Butterworth-Heinemann.

Garrick, B.J. 1988. The approach to risk analysis in three industries: nuclear power, space systems, and chemical process. Reliability Engineering and System Safety 23(3): 195–205.

Garrick, B.J. 1989a. Risk assessment practices in the space industry: the move toward quantification. Risk Analysis 9(1): 1–7.

Garrick, B.J. 1989b. Lessons learned from 21 nuclear plant probabilistic risk assessments. Nuclear Technology 84: 319–330.

Garrick, B.J. 1999. Risk Assessment Methodologies Applicable to Marine Systems. Washington, D.C.: U.S. Coast Guard.

Garrick, B.J., and R.F. Christie. 2001. Probabilistic risk assessment practices in the United States of America for nuclear power plants. Safety Science (Version 7): 25–49.

Garrick, B.J., and S. Kaplan. 1995. A Risk-Based Approach to the Evaluation and Application of Disarmament Strategies and Technologies, Summary Report. Prepared for North Atlantic Treaty Organization, Advanced Research Workshop, May 19-23, 1995. Visegrad, Hungary.

Garrick, B.J., and S. Kaplan. 1999. The Security and Safety Risk of the Space Shuttle Vehicle. Prepared for United Space Alliance. Houston, Tx: United Space Alliance.

Haimes, Y. 1998. Risk Modeling, Assessment, and Management. New York: John Wiley & Sons, Inc.

Helton, J.C., D.R. Anderson, G. Basabilvazo, H.N. Jow, and M.G. Marietta. 2000. Conceptual structure of the 1996 performance assessment for the Waste Isolation Pilot Plant. Reliability Engineering and System Safety 69: 151–165.

Heuer (Jr.), R.J. 1999. Psychology of Intelligence Analysis. Washington, D.C.: Center for the Study of Intelligence, Office of the Director of Central Intelligence.

Hora, S.C., D. von Winterfeldt, and K.M. Trauth. 1991. Expert Judgment on Inadvertent Human Intrusion into the Waste Isolation Plant. SAND90-3063. Albuquerque, N.M.: Sandia National Laboratories.

Howard, R.A., and J.E. Matheson. 1989. Readings on the Principles and Applications of Decision Analysis. Menlo Park, Calif.: Strategic Decisions Group.

Johnson, D.H., et al. 1997. B-52H Electrical Environments Modeling and Probabilistic Risk Assessment Inputs. DSWA-TR-96-88. Alexandria, Va.: Defense Special Weapons Agency.

Kaplan, S. 1992. Expert information versus expert opinions: another approach to the problem of eliciting/combining/using expert judgment in PRA. Reliability Engineering and System Safety 35: 61–72.

Kaplan, S., and B.J. Garrick. 1981. On the quantitative definition of risk. Risk Analysis 1(1): 11–27.

Lewis, H.W. 1990. Technological Risk. New York, London. W.W. Norton & Company.

Lawrence Livermore National Laboratory). 1995. Recommendations for Probabilistic Seismic Hazard Analysis: Guidance on Uncertainty and the Use of Experts. Sponsored by the U.S. Department of Energy, U.S. Nuclear Regulatory Commission, and Electric Power Research Institute. UCRL-ID-122160 (2 Vols.) Livermore, Ca.: Lawrence Livermore National Laboratory.

Molak, V., ed. 1997. Fundamentals of Risk Analysis and Risk Management. Cincinnatti, Oh. Lewis Publishers.

NASA (National Aeronautics and Space Administration). 2002. Probabilistic Risk Assessment Procedures Guide for NASA Managers and Practitioners. Washington, D.C.: Office of Safety and Mission Assurance, NASA Headquarters.

NRC (National Research Council). 1988. Post-Challenger Evaluation of Space Shuttle Risk Assessment and Management. Washington, D.C.: National Academy Press.

NRC. 1989. Improving Risk Communication. Washington, D.C.: National Academy Press.

NRC. 1994a. Building Consensus through Risk Assessment and Management of the Department of Energy's Environmental Remediation Program. Washington, D.C.: National Academy Press.

NRC. 1994b. Recommendations for the Disposal of Chemical Agents and Munitions. Washington, D.C.: National Academy Press.

NRC. 1996a. Understanding Risk: Informing Decisions in a Democratic Society. Washington, D.C.: National Academy Press.

NRC. 1996b. The Waste Isolation Pilot Plant: A Potential Solution for the Disposal of Transuranic Waste. Washington, D.C.: National Academy Press.

NRC. 1999. An End State Methodology for Identifying Technology Needs for Environmental Management. Washington, D.C.: National Academy Press.

NRC. 2002. Making the Nation Safer: The Role of Science and Technology in Countering Terrorism. Washington, D.C.: The National Academies Press.

Raiffa, H. 1996. Decision Analysis. Columbus, Oh. McGraw-Hill Primis Custom Publishing.

USNRC (U.S. Nuclear Regulatory Commission). 1996. Branch Technical Position on the Use of Expert Elicitation in the High-Level Radioactive Waste Program. NUREG-1563. Washington, D.C.: U.S. Nuclear Regulatory Commission.

USNRC. 1999. White Paper on Risk-Informed and Performance-Based Regulation. SECY-98-144. Washington, D.C.: U.S. Nuclear Regulatory Commission.

Woods, K.S. 1997. Quantitative Risk Assessment Final Report. Washington, D.C.: National Aeronautics and Space Administration.

CHAPTER 3

THE FOUNDATION FOR QUANTITATIVE RISK ASSESSMENT

This chapter reviews a few of the basic concepts of quantitative risk assessment (QRA). The goal is to give the reader an understanding of what is involved in assessing the risk of terrorist attacks using well known quantitative techniques in preparation for the sample application provided in Chapter 4.

BASIC REQUIREMENTS

QRA must meet the following basic requirements to support "making the right decisions" on ways to combat terrorism:

- Quantification of uncertainty must be an inherent feature of the methodology.
- The methodology must be general, that is, it should be applicable to all types of risk to maximize the use of the existing experience base in the risk sciences and to accommodate new methods as they evolve.
- The methodology must be supported with an experience base that demonstrates proof of concept.
- The methodology must be specialized for application to those threats deemed most serious in terms of catastrophic consequences. Existing *qualitative* methods of risk assessment should be used to screen out attack scenarios where there is ample evidence that such screening is appropriate.

COMBATING TERRORISM THROUGH THE QUANTITATIVE RISK ASSESSMENT PROCESS

The remainder of this chapter focuses on quantitative risk assessment as a methodology. The study group believes that a systematic, quantitative risk assessment will help answer questions about how to manage the risk of terrorism—questions such as: (1) what are the threats and vulnerabilities of greatest importance; (2) what are the risk-contributing factors and how do they rank in importance; and (3) what actions will have the biggest payoff in terms of risk reduction for the amount of resources invested? Quantitative risk assessment requires the development of a set of scenarios describing what constitutes successful operation of a system and how the system can fail or be made to fail, catastrophically or otherwise.

Guiding Principles

Guiding principles for scenario-based risk assessment applied to combating terrorism are listed below:

- The quantitative expression of risk should be in the form of a structured set of scenarios, each having a corresponding likelihood and consequence.
- The set of scenarios must be complete in the sense that all of the important contributors to risk are included.
- The scenarios must be quantified in terms of clearly defined risk measures, must be realistic, and must incorporate uncertainties.
- Each scenario should depict a terrorist attack in the form of a sequence of events, starting with the initiating event that upsets an otherwise successful operation or system and proceeding through a series of subsequent events to the end-state (i.e., the consequences of the attack). The initiating event must be based on a comprehensive threat assessment.
- Each scenario must accommodate combined events, including primary and diversionary events.
- The end-states must reflect initial, cascading, and collateral consequences (or levels of damage).
- Individual events and aggregated event uncertainties must be quantified on the basis of the evidence.
- The results must be ranked as to their contribution to risk in order of importance and must be presented in a way that supports decision-making.

These principles, and the triplet definition of risk, are the basis for the methodology. It is important to note that except in the area of uncertainty analysis the concepts and ideas have all been tried and tested in extremely diverse applications, from sabotage and security to accidents; from transportation of hazardous materials to processing and manufacturing plants; from offshore oil platforms to the risk of importing exotic animals; and from food safety to nuclear power plants. The difference between applications is not in the fundamental concept, but the boundary conditions of specific applications, and especially in the way the input data are prepared.

Implementation of the Principles

Adherence to these principles is achieved through the following six-step process:

1. Define the system being analyzed in terms of what constitutes normal operation and points of vulnerability to serve as a baseline reference point.
2. Identify and characterize the "sources of danger," that is, the hazards (e.g., stored energy, toxic substances, hazardous materials, acts of nature, sabotage, terrorism, equipment failure, combinations of each, etc.).
3. Develop terrorist attack scenarios to establish levels of damage and consequences.
4. Adopt risk metrics that reflect the likelihoods of different attack scenarios in terms of target and collateral damage and quantify the scenarios based on the totality of relevant evidence.

5. Assemble the scenarios according to damage levels, and cast the results into the appropriate risk curves and risk priorities.
6. Interpret the results to guide the risk-management process.

Fundamental QRA Concepts

The fundamental concepts of quantitative risk assessment include an understanding of basic definitions, the quantification of uncertainty, the structuring and quantification of scenarios, and the assembly of results in a form that supports effective risk management.

Defining Risk

In general, risk assessments satisfy the "triplets definition of risk," *R*, which can be expressed as follows (Kaplan and Garrick, 1981),

$$R = \{<S_i, L_i, X_i>\}_c$$

where S_i denotes risk scenario i, L_i denotes the likelihood of that scenario, and X_i denotes the consequences of that scenario. The angle brackets < > enclose the triplets, the curly brackets mean "a set of," and the subscript *c* denotes complete, meaning that all of the important scenarios are included in the set. The body of methods used to identify the scenarios (S_i) constitutes what is called the theory of scenario structuring (Kaplan, et al., 2001).

In accordance with the triplet definition, quantification of risk entails answering three questions: (1) what can go wrong? (2) how likely is it? and (3) what are the consequences? When applied to the risk of terrorism, the question of what can go wrong has a different spin. The question becomes: how can a terrorist deliberately make something happen to achieve a desired outcome? The principle is the same, but the perspective is very different (Box 3-1).

Box 3-1. A Simple Example Using the Triplet Definition of Risk

The triplet definition of risk says that a risk assessment can be thought of as a structured set of scenarios (*S*), likelihoods (*L*), and consequences (*X*). For example, suppose the problem was determining the risk of a team of hikers taking on a challenging hike in the well-known primitive area of Idaho. The first thing would be to establish what has to happen for the hike to be successful. We call this the "success" scenario and it is the basis for knowing what is considered a departure from success. The second thing would be to ask what could go wrong. Of course, there are many possibilities, but suppose we focus on the risk of a catastrophic event such as serious injuries, fatalities, or getting completely lost. We do this because it's the catastrophic events that we really don't want to happen, and it keeps us from getting caught up in the details of events that may be annoying and discomfiting, but not relevant to the hikers experiencing a catastrophe.

Asking the question "what can go wrong" that would be catastrophic conjures up all kinds of possibilities, such as being attacked by a wild animal, encountering an unexpected severe storm, being attacked by bandits, experiencing an earthquake or a forest fire, being crushed by a landslide, having a bad accident, going berserk, etc. The idea is to develop a complete set of the most important threat scenarios. The next thing is to evaluate the consequences of the scenarios and screen out the ones that do not result in catastrophic consequences. Finally, we consider whatever evidence we can find on the likelihood of each scenario. Evidence can be in the form of experience with similar hikes, the degree of difficulty of the hike, the experience and behavior of the members of the hiking team, and the susceptibility of the region to natural events, such as floods, storms, earthquakes, and fires. Based on the evidence, we then assign a chance factor, or probability, to each scenario. We now have a structured set of scenarios, their likelihoods, and consequences. We can now make informed decisions about the hike. Of course, we could get very sophisticated and manipulate the scenarios into different forms, such as totaling the risk of all of the scenarios and casting the results into risk curves and tables. But for our purposes, we can just make a decision on the basis of the scenarios, likelihoods, and consequences.

Quantifying Uncertainty and Bayes Theorem

The central feature of quantitative risk assessment is making uncertainty an inherent part of the analysis. Uncertainty refers to the parameters that are used to measure risk and how these parameters represent uncertainties in information and modeling. One theory fundamental to quantifying the uncertainty in risk is Bayes Theorem, "a striking advance in statistics by demonstrating how to make better-informed decisions by mathematically blending new information into old information" (Bernstein, 1996). Simply put, Bayes Theorem is the fundamental, logical principle governing the process of inferential reasoning. The theorem answers the question: "How does the probability *p(h)*, of a given hypothesis, *h*, change when we obtain a new piece of evidence, *E*?" The answer is very simply derived as follows. Let *p(h|E)* denote the new probability of *h*, given that we now have the evidence E, and let $p(h \wedge E)$ denote the probability that both *h* and *E* are true. Then,

$$p(h \wedge E) = p(E)p(h|E) \qquad (1)$$

This equation simply says that the probability of both *h* and *E* being true is equal to the probability that *E* is true times the probability that *h* is true, given the evidence *E*. In the same way,

$$p(h \wedge E) = p(h)p(E|h) \qquad (2)$$

Setting the two right-hand sides of (1) and (2) equal to each other, and dividing by *p(E)* gives the following result.

$$p(h/E) = p(h)\,\frac{p(E|h}{p(E)} \qquad (3)$$

Equation (3) is Bayes Theorem. It tells us how the probability of a hypothesis *h* changes when we learn a new piece of evidence *E* (Box 3-2).

Box 3-2. Bayes Theorem: An Example

In Box 3-1, the hiking team is aware of the possibility of encountering a bear along the trail. Based on statistical averages published by the forest ranger's office, they assign a 50 percent probability to the hypothesis that they will indeed have such an encounter. In Bayesian language this is called the prior probability. But the hikers also know that the frequency of bear encounters may depend on details not evident in the information, such as the date of departure, the route chosen, weather conditions, etc. The hikers would like to take these details into account.

They go to the forest ranger and tell him the date and route of the planned hike. The ranger provides evidence showing that bears would still be hibernating at that time of the year, so there would be very little risk of an encounter. The hikers update their estimate based on this new evidence. This updating is exactly what Bayes Theorem does. Let us suppose that the result of this updating is to reduce the probability of an encounter by a factor of 10. Thus, all of the evidence taken together yields a 5 percent chance of encountering a bear, instead of the 50 percent chance estimated based on more global data. The hikers decide that with only a 5 percent chance of encountering a bear, the risk is worth taking, and they proceed with their plans.

It is conceivable that this equation could be programmed into a "surveillance computer" in such a way that when items of relevant evidence become known, the computer will call attention to the fact that the probability of a terrorist threat has increased. Given this forewarning, appropriate action can be taken (Box 3-3).

Box 3-3. Combating a Terrorist Attack During the Planning Stage

Developing a logic diagram for the preparation of a terrorist attack offers an opportunity to diagnose and determine if such preparations are being made, thus allowing preemptive actions to be taken. The attack plan might include specific activities, such as moving people and money around, purchasing equipment and supplies, establishing communications, reconnoitering, and so forth, all activities necessary for acquiring the resources or capabilities to implement the plan. This suggests that, for selected critical infrastructures, a surveillance system could be set up to identify such activities, which would then provide evidence from which we might infer that a terrorist attack plan is being carried out.

These inferences would be made, using Bayes Theorem, by computers monitoring the surveillance system. Using all the evidence from the surveillance system, as well as all other relevant information, the probability that preparations for a terrorist attack are under way could be calculated, as well as insights into the status, stage, and nature of the preparations. These calculations would be done on line, automatically and continuously.

The relationship between the surveillance process and the logic modeling discussed in this report is illustrated in the figure below. If the probability that an attack is under way increases significantly, the system calls it to the attention of the authorities so they can take defensive action.

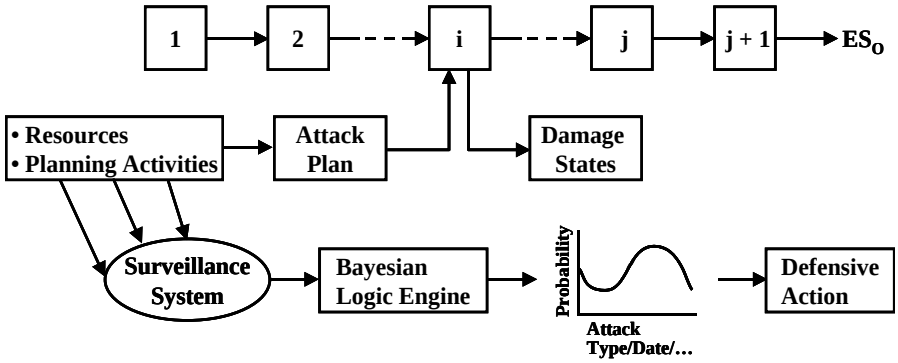

It should be noted that the "computer" is not necessarily an electronic machine. It could be a human being surveying the incoming evidence. Human brains also operate according to Bayes Theorem, but they are not as fast or as reliable. It is best to combine the strengths of both. Thus, the computer, programmed with Bayes Theorem, helps humans to "connect the dots" between different sources of information and take the right actions. Bayesian methods, including uncertainty analysis, have already been be used to model terrorist attacks (Paté-Cornell, 2001; Paté-Cornell and Guikema, 2002).

Structuring the Scenarios

Scenario structuring encompasses the methods, algorithms, and insights needed to identify and portray the risk scenarios *(S_i).* It is convenient to structure terrorist attack scenarios from the point of view of the system that is attacked. An example of an initial point of attack is terrorists taking over an airplane loaded with fuel to use as a weapon. Threat assessment is then the task of quantifying "initiating events" that disrupt an otherwise normally operating system.

The first step in the process of structuring scenarios is to develop a diagram describing the "success scenario" *(S_0)* that leads to a successful end-state *(ES_0)* or normal operating procedures for the system without the intervention of a terrorist event. In other words, the success scenario describes the functioning of the system when it is working as planned. It usually, but not necessarily, has a linear structure of events as depicted in Figure 3-1 where the *i* and *j* simply represent any number of components that could become the initial target of the terrorists. An ordinary risk assessment asks what can go wrong with each part of the diagram,

especially those parts considered most vulnerable to the risk being considered. The answers to this question are called initiating events (*IE*) because they initiate the risk scenarios *(S_i).*

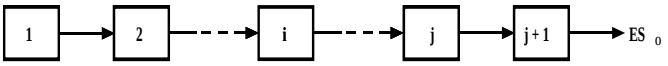

Figure 3-1. Diagram of a Success Scenario

Given that an *IE* has occurred, an event tree then emerges (Box 3-4, Figure 3-2). Each path through this tree represents a scenario S_i and ends up at an end state *(ES_i),* initiated by a terrorist event.

Box 3-4. Use Fault Tree and Event Tree Methodology to Link Threat and Vulnerability Assessment

The most common logic diagrams used by practitioners of quantitative risk assessment are event trees and fault trees. The two complement each other. An event tree starts with an initiating event and proceeds to identify succeeding events, including branches that eventually terminate into possibly undesirable consequences. An event tree, therefore, is a cause-and-effect representation of logic.

A fault tree starts with the end-state or undesired consequence and attempts to determine all of the contributing system states. Therefore, fault trees are effect-and-cause representations of logic. An event tree is developed by inductive reasoning while a fault tree is based on deductive reasoning. A key difference in the two representations is that a fault tree is only in "failure space," and an event tree includes both "failure and success space." The choice between the two is a matter of circumstances and preference, and they are often used in combination; the event tree provides the basic scenario space of events and branch points, and the fault tree is used to quantify the "split fractions" at the branch points.

The fault tree is useful for investigating activities to achieve a "desired outcome," such as how to successfully attack the electric grid system; event trees examine the possible outcomes of different attacks (i.e., "initiating events" that upset an otherwise normally operating system.) Figure 3-3 illustrates how threat assessment and vulnerability assessment are linked to create a representation of the integrated attack scenario.

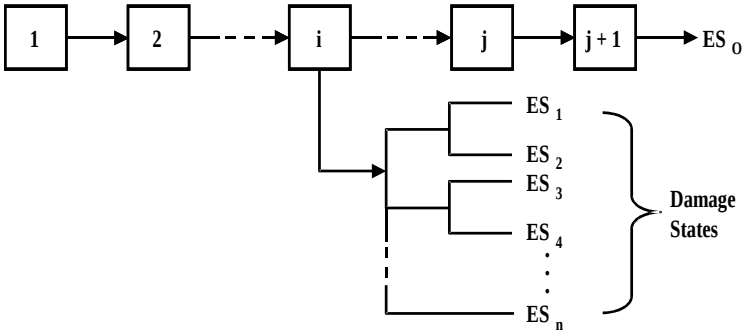

Figure 3-2. An Event Tree Showing Scenarios Emerging from an Initiating Event

A terrorist planning an attack and the risk analyst must both think in terms of scenarios or sequences of events. This becomes the core of the threat assessment. The terrorist asks "how he or she can make something go wrong". In other words, (s)he asks how (s)he might introduce an event into a normal operation that would facilitate achieving the ultimate goal, the end-state of the attack scenario. The task of developing meaningful "initiating events" and assessing their likelihoods is considered by many to be the most difficult challenge in quantifying the risk of a terrorist attack. This is the threat assessment part of a terrorist attack risk assessment.

The terrorist must create a plan based on his capabilities and intentions. Required resources can include people, money, information, delivery systems, and weapons. The terrorist's plan (intentions and capabilities) can be represented as an incoming fault tree to the initiating or triggering event of the attack, as illustrated in Figure 3-3. This is often described in the QRA field as a master logic diagram for developing initiating events. It has a fault-tree-like structure that displays the events and resources used to "initiate" an attack.

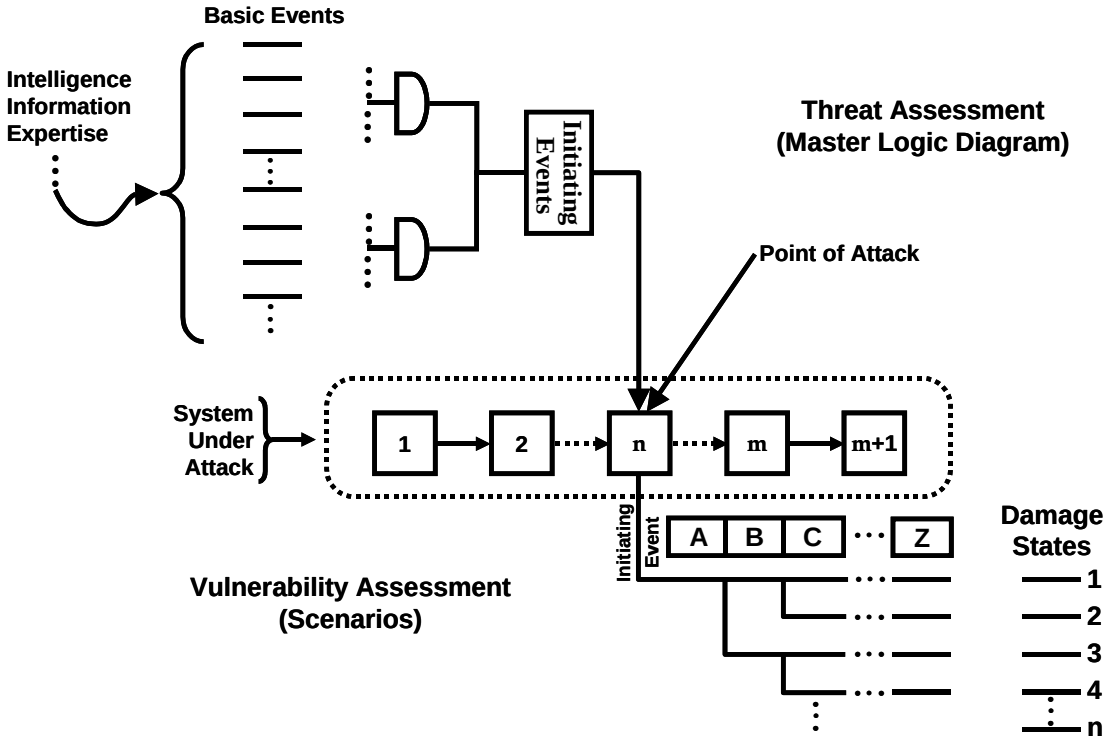

Figure 3-3. The Concept of an Integrated Threat and Vulnerability Risk Assessment

The master logic diagram is based on input from intelligence and terrorism experts and their knowledge of activities and preparations the terrorist must make to launch an attack at different points of the S_0 scenario, that is, the scenario that represents how the system normally works before it is attacked. Figure 3-3 illustrates how threat assessment and vulnerability assessment are linked to create a representation of the integrated attack scenario. The pinch point between threat assessment and vulnerability assessment is the initiating event.

Using risk assessment techniques to analyze terrorist attacks may appear to be much more problematic than, for example, analyzing the risk of a fixed system, such as a manufacturing plant of hazardous materials. However, the systems that can come under attack from a terrorist are also fixed and well defined. The differences are in the nature of the threats rather than in the systems. Threats can be accidents and external events, such as fires, severe storms, earthquakes, aircraft impacts, sabotage, and, of course, terrorism. The accident part of risk assessment is tied to the design and operations of the facility, which are usually well defined. But once the risk assessment seeks to be complete in the sense of considering "external" threats,

the differences between the two applications diminish. In fact, most large scope quantitative risk assessments now consider external as well as internal events, although only recently has the threat of a terrorist attack been seriously considered.

Quantifying the threat of a terrorist attack represents a challenge, primarily because it is a new threat to be considered, rather than because of any intrinsic characteristic of threat assessment; the principles of the assessment are the same whether it involves a rare and severe storm or a terrorist attack—a systematic consideration of the evidence and an involvement of the appropriate experts. It may turn out that the uncertainties make it difficult to make a decision, but that is not a computational problem of the risk assessment methodology.

The Concept of "Likelihood" and the "Probability of Frequency" Framework

To quantify the likelihood of attack scenarios, it is first necessary to define the concept of likelihood. So far, we have purposely used the term "likelihood" as a general, intuitive expression in the triplet definition of risk. Now we describe three explicit and quantitative interpretations of likelihood. These are "frequency," "probability," and "probability of frequency."

Frequency. If the scenario is recurrent, that is, if it happens repeatedly, then the question "how frequently" can be asked, and the answer can be expressed in occurrences per day, per year, per trial, per demand, etc.

Probability. If the scenario is not recurrent (i.e., if it happens either once or not at all), then its likelihood can be quantified in terms of "probability." "Probability", in our usage, is synonymous with "credibility." Thus "probability" is the degree of credibility of the hypothesis in question, based on the totality of relevant evidence available.

Probability of Frequency. If the scenario is recurrent, and therefore has a frequency, but the numerical value of that frequency is not fully known, and if there is some evidence relevant to that numerical value, then Bayes Theorem (the fundamental principle governing the process of making inference from evidence) can be used to develop a probability curve over the frequency axis. This approach has been widely used in the risk assessment of engineered systems. This "probability of frequency" interpretation of likelihood is the most informative, and thus is the preferred way of capturing and quantifying the state of knowledge about the likelihood of a defined scenario.

Having defined what we mean by likelihood, we can explain our usage of "probability." What some call the "subjectivist" view of probability is best expressed by the physicist E.T. Jaynes (2003): "A probability assignment is 'subjective' in the sense that it describes a state of knowledge rather than any property of the 'real' world, but is 'objective' in the sense that it is independent of the personality of the user; two beings faced with the same background of knowledge must assign the same probabilities." The central idea of Jaynes is to bypass opinions and seek out the underlying evidence for the opinions which therefore becomes more objective rather than subjective.

We agree wholeheartedly with Jaynes' statement and, in our usage, go yet a step further. We define "probability" as synonymous with "credibility." We can thus speak, and think, in terms of the "credibility" of a hypothesis based on all the evidence available! "Credibility" is thus a positive number ranging from zero to one, and it obeys Bayes Theorem. Thus, if we write *p(h|E)* to denote the credibility of hypothesis *h*, given *E*, then

$$p(h/E) = p(h)\,\frac{p(E|h}{p(E)}$$

which is Bayes Theorem, and which tells us how the credibility of hypothesis *h* changes when new evidence, *E*, occurs. It does that without overt reference to a "user" or "sentient beings"—it is completely objective as it is only evidence based, not opinion or personality based.

The debate between the so-called subjectivists and the frequentists, sometimes referred to as the Bayesians and the classical statisticians is legendary and has been going on for over 200 years. This debate has been the subject of textbooks and scientific articles on probability since the time of LaPlace and Bayes, a few of which are referenced (de Finetti, 1974), (Apostolakis, 1990), (Lindley, 1985).

Quantifying Initiating Events

Before the risk scenarios can be quantified, initiating events of the scenarios must be described and quantified. A QRA-based threat (initiating event) assessment involves the following tasks:

1. Gaining access to all the evidence of a terrorism threat. That evidence will be in the form of intelligence information, data, and analyses of past terrorist attacks.
2. Assembling appropriate experts to interpret evidence and screen out threats that do not meet the established criteria. The goal is to target threats with catastrophic consequences for which there is intelligence of an imminent threat.
3. Constructing a model of the threats in the form of a fault tree that provides the logic between the initiating event (the threat) and the basic inputs as depicted in Figure 3-3.
4. Reducing the observations noted in Step 1 to obtain the basic events required in Step 3.
5. Exercising the threat model to quantify the selected set of initiating events.

Tasks 4 and 5 are simply the application of the principles of the scientific process of reducing observations to numbers. A deductive logic model, that is, a fault tree or master logic diagram, is developed in Task 3 for each initiating event (see Figure 3-3) of the screened set from Task 2. The structure of the logic model is to deduce from the selected set of hypothetical initiating events the intervening events down to the point of the *intentions* of the terrorist, that is, the *decision to launch* an attack. The intention of the terrorist or the decision to launch an attack is the "basic event" of the threat logic diagram. The intervening events of the logic diagram are representations of the planning, training, logistics, resources, activities, and capabilities of the terrorists. The tools for developing the basic events as depicted in Figure 3-3 are rooted in expert elicitation, evidence provided by the experts, and inferential reasoning.

Quantifying the Scenarios

The actual quantification of the risk scenarios is done with the aid of an event tree (Figure 3-4). An event tree is a diagram that traces the response of a system to an initiating event, such as a terrorist attack, to different possible end points or outcomes (consequences). A single path through the event tree is the scenario. The event tree displays the systems, equipment, human actions, procedures, etc., that can impact the consequences of an initiating event depending on the success or failure of intervening actions. In Figure 3-4 boxes with the letters A, B, C, and D represent these intervening actions. The general convention is that if the action is successful, the scenario is mitigated. If the action is unsuccessful, then the effect of the initiating event continues as a downward line from the branch point in Figure 3-4. An example of an action that could mitigate the hijacking of a commercial airliner to use it as a weapon to crash into a football stadium would be a remote takeover of the airplane by ground control or shooting it down.

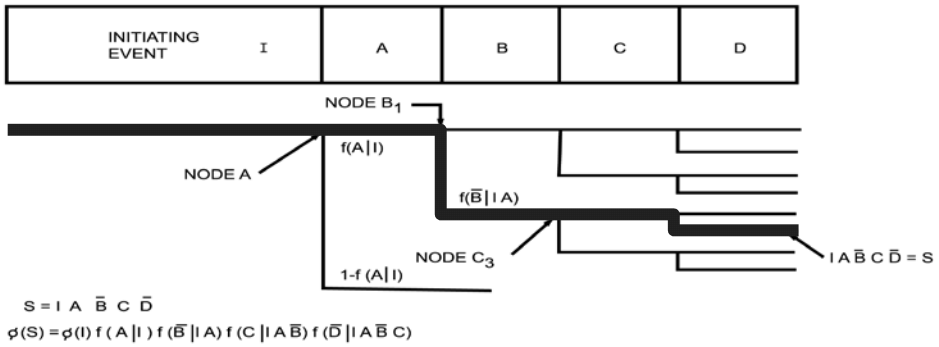

Figure 3-4. Quantification of a Scenario Using an Event Tree

Each branch point in the event tree has a probability associated with it. It should be noted that the diagram shown in Figure 3-4 shows only two branches (e.g., success or failure). However, an event tree can also have multiple branches to account for different degrees of degradation of a system. These branch points have associated "split fractions" that must be quantified based on the available evidence. The process involves writing an equation for each scenario of interest. For example, the path through the event tree that has been highlighted in Figure 3-4 could be a scenario that we wish to quantify. The first step is to write a Boolean equation, an algebraic expression, for the highlighted path. If we denote the scenario by the letter *S*, we have the following equation,

$$S = IA\bar{B}C\bar{D}$$

where the bars over the letters indicate that the event in the box did not perform its intended function. The next step is to convert the Boolean equation into a numerical calculation of the frequency of the scenario. Letting φ stand for frequency and adopting the split fraction notation, *f(…)*, of Figure 3-4, gives the following equation for calculating the frequency of the highlighted scenario,

$$\varphi(S) = \varphi(I)f(A/I)f(\bar{B}/IA)f(C/IA\bar{B})\,f(\bar{D}/IA\bar{B}C)$$

The remaining step is to embed the frequencies into appropriate probability distributions to communicate their uncertainties. This is done using Bayes Theorem to process the elemental parameters (Figure 3-5). The "probability of frequency" of the individual scenarios is obtained by convoluting the elemental parameters in accordance with the above equation.

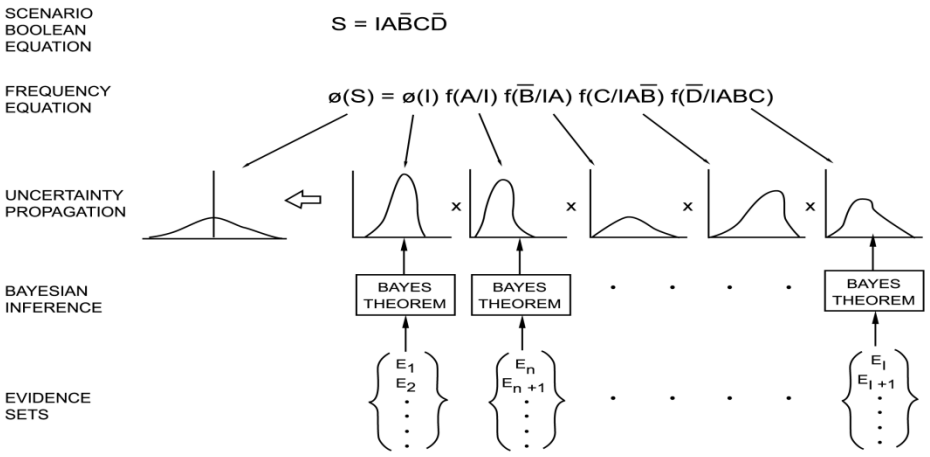

Figure 3-5. Bayes Theorem Used to Process Parameters

Assembling the Results

Once the scenarios have been quantified, the results take the form of the graph in Figure 3-6. Each scenario has a probability-of-frequency curve quantifying its likelihood of occurrence. Figure 3-6 shows the curve for a single scenario or a set of scenarios leading to a single consequence. Showing different levels of damage, such as the risk of varying injuries or fatalities, requires a different type of presentation. The most common form is the classical risk curve, also known as the frequency-of-exceedance curve, or the even more esoteric label, the complementary-cumulative-distribution-function. This curve is constructed by ordering the scenarios by increasing levels of damage and cumulating the probabilities from the bottom up in the ordered set against the different damage levels. Plotting the results on log-log paper generates curves, as shown in Figure 3-7.

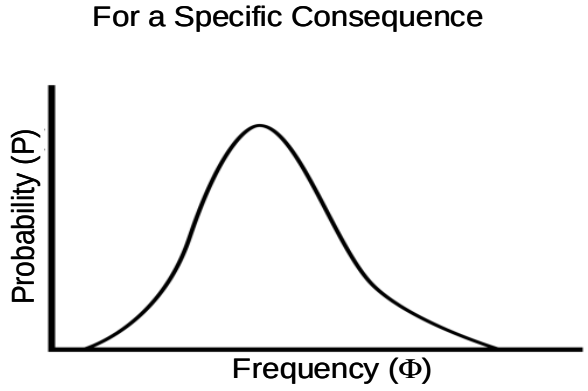

Figure 3-6. Probability-of-Frequency Curve

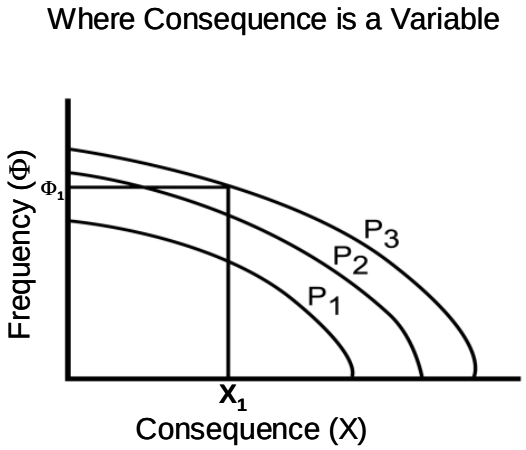

Figure 3-7. Risk Curve for Varying Consequences

To illustrate how to read Figure 3-7, suppose P_3 has the value of 0.95, that is a probability of 0.95, and suppose we want to know the risk of an X_1 consequence at the 95 percent confidence level. According to the figure, we are 95 percent confident that the frequency of an X_1 consequence or greater is Φ_1. The family of curves (usually called percentiles) can include as many curves as necessary. The ones most often selected in practice are the 5th, 50th, and 95th percentiles. A popular fourth choice is the mean.

Although risk assessment results such as those illustrated in Figures 3-6 and 3-7 can be beneficial in providing a perspective on the actual risks and in establishing priorities for threats, targets, and vulnerabilities, they are not the most important output of the risk assessment. The most important output is the revelation of the dominant contributors to the risk, which must be identified for effective risk management. The contributors are buried in the results assembled to generate the curves in Figures 3-6 and 3-7. Most risk assessment software packages contain algorithms for ranking the importance of contributors to a risk metric.

CONCLUSION AND RECOMMENDATION

Conclusion. Risk assessment based on quantitative methods and decision analysis has contributed to safer, more environmentally acceptable products, services, and systems in several industries.

Recommendation. Government and industry should increase the use of quantitative risk assessment to support decisions for combating terrorism. The U.S. Department of Homeland Security should issue policy guidelines for implementing a quantitative risk assessment process based on scientific principles that integrates threats and vulnerabilities, clearly links the decision options with supporting evidence, and displays the characteristics of risks, benefits, and costs.

REFERENCES

Apostolakis, G. 1990. The Concept of Probability in Safety Assessments of Technological Systems. Science, 250(1990) 1359-1364.

Bernstein, P.L. 1996. Against the Gods: The Remarkable Story of Risk. New York: John Wiley & Sons.

de Finetti, B. 1974. Theory of Probability, Vols. 1 and 2, New York, 1974

Jaynes, E.T. 2003. Probability Theory: The Logic of Science. Cambridge, U.K.: Cambridge University Press.

Kaplan, S., and B.J. Garrick. 1981. On the quantitative definition of risk. Risk Analysis 1(1): 11-27.

Kaplan, S., Y.Y. Haimes, and B.J. Garrick. 2001. Fitting hierarchical holographic modeling into the theory of scenario structuring and a resulting refinement to the quantitative definition of risk. Risk Analysis 21(5): 807-819.

Lindley, D.V. 1985. Making Decisions. Second Edition. Wiley, London.

Paté-Cornell, E. 2001. Fusion of intelligence information: a Bayesian Approach. *Risk Analysis* 22(3): 445-454.

Paté-Cornell, E., and S. Guikema. 2002. Probabilistic modeling of terrorist threats: a systems analysis approach to setting priorities among countermeasures. *Military Operations Research* 7(4): 5-20.

CHAPTER 4

ASSESSING THREATS AND VULNERABILITIES: A SAMPLE APPLICATION[3]

The purpose of this chapter is to provide an example that demonstrates some of the key features of the methodology described in Chapters 2 and 3. Because it is not possible to include or summarize an actual risk assessment of a terrorist attack (for reasons of security, resource limitations, availability, etc.), a hypothetical but realistic example is used to show how the major steps of a quantitative risk assessment (QRA) are implemented.

Considering how pervasive electricity is to the functioning of a society, the study group decided to present a risk assessment of a hypothetical electrical grid. The electricity infrastructure is critical to the nation's well being and is currently under study for its vulnerability to terrorist attacks (Amin, 2002; EPRI, 2002). The sector is already responding to the threat of possible terrorist attacks through risk-management practices (NAERC, 2002). Risk management is an integral part of the electricity sector's definition of "critical infrastructure protection," and is defined as "safeguarding the essential components of the electric infrastructure against physical and electronic threats in a manner consistent with appropriate risk management, with both industry and industry-government partnerships, while sustaining public confidence in the electricity sector."

The example involves a risk assessment of a combined cyberattack (Alvey, 2002) and physical attack on a hypothetical electric power grid. Even though the scope of the assessment is limited, it demonstrates how risk assessment can help in decision-making. The threat assessment part of the example is more limited than the vulnerability assessment, primarily because of the lack of resources to search out threat information. *The threat assessment of the cyberattack includes some of the steps leading up to the attack (the initiating event for the vulnerability assessment) but does not assess or speculate on the terrorists' decision to initiate the attack. The assessment of the physical attack portion simply assumes that the attack takes place.* The vulnerability assessment takes the consequences to the point of inflicting damage to the grid. However, the analysis does not include health and safety effects or long-term cascading economic and environmental impacts that might result. The authors have attempted to provide enough detail to convey the ideas of the methodology without resorting to overly technical jargon. For the principal audience of this report (policy makers, decision makers, etc.), there may be too much detail. For the technical community, there may not be enough detail. The authors suggest that policy makers concentrate on the first few and last few pages of the chapter, which cover the essentials; more technically inclined readers may want to look into the references in Chapters 2 and 3 for details on the analytical steps of the proposed risk assessment process.

[3] John W. Stetkar, an independent consultant to the electric power industry, provided support to the study group for this chapter.

The example shows how vital systems can be analyzed to expose their vulnerabilities and provide a basis for taking corrective actions either to avert or mitigate the consequences of a terrorist attack. The risk assessment of the sample electrical grid leads to specific recommendations, derived from the supporting evidence, which could not have been easily deduced, or supported, without this formal approach.

The example follows the six-step process introduced in Chapter 3.

1. Define the system being analyzed in terms of what constitutes normal operation and points of vulnerability to serve as a baseline reference point.
2. Identify and characterize the "sources of danger," that is, the hazards (e.g., stored energy, toxic substances, hazardous materials, acts of nature, sabotage, terrorism, equipment failure, combinations of each, etc.
3. Develop terrorist attack scenarios to establish levels of damage and consequences.
4. Adopt risk metrics that reflect the likelihoods of different attack scenarios in terms of target and collateral damage and quantify the scenarios based on the totality of relevant evidence.
5. Assemble the scenarios according to damage levels, and cast the results into the appropriate risk curves and risk priorities.
6. Interpret the results to guide the risk-management process.

Step 1. Define the system being analyzed in terms of what constitutes normal operation and points of vulnerability to serve as a baseline reference point.

DEFINING THE SYSTEM

The purpose of the first step is to understand how the system works so departures from normal, successful operation can be easily identified. Once the system is understood, vulnerabilities that require special analysis can be identified.

In an increasingly interconnected world, technology-based systems and networks are becoming more and more interdependent. An attack on one system can have far-reaching, cascading effects on other systems and on society as a whole. The system in this example is a hypothetical portion of a national electric power grid, which is tightly linked to other vital systems and, therefore, is an attractive target for terrorists. The consequences of an attack on a national electric power grid that leads to long term-outages, say greater than 48 hours, could cascade into major disruptions in transportation, communications, sanitation, food supplies, water supplies, and other systems.

The Region

Figure 4-1 represents a major region in the national electric power grid; each network corresponds to a large metropolitan area, such as New York City, Philadelphia, Boston, etc. Networks are interconnected to form a regional grid (such as the northeast corridor or the

western states). In Figure 4-1, Network 1 is interconnected with four neighboring networks through ties T12, T13, T14, and T15. These "interties" (pronounced "inter-ties") form the transmission system and are typically extra-high voltage (EHV) transmission lines that provide the major pathways for power flow throughout the region and between cities. Regional grid operations are typically coordinated through established protocols designed to ensure economical transfers of power through the interties and to prevent failures from cascading and causing widespread disruptions in power (such as those that occurred August 2003.)

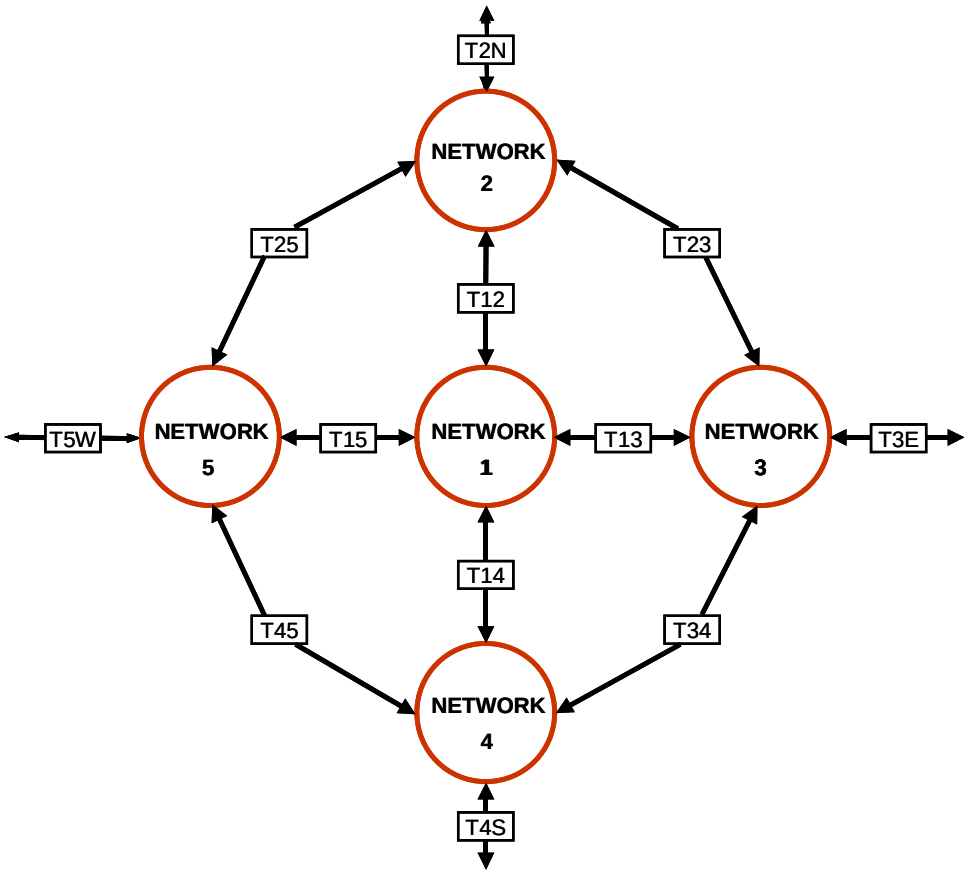

Figure 4-1. Sample Regional Grid

Figure 4-1 shows that external power can be routed to Network 1 through several parallel interties. In some parts of the country, the available interconnections are limited; well known examples include the north-south ties through the Western Interconnection, ties from the southern power pools to the Electric Reliability Council of Texas Interconnection, and limited ties to Florida through the Eastern Interconnection. Because regional-specific features must be taken into account, risk assessments cannot be performed generically.

In addition, the U.S. Department of Energy has identified several transmission bottlenecks at various interties throughout the U.S. electrical grid (DOE, 2002). "Bottlenecks" occur at points where major tie lines are frequently loaded to a large fraction of their available capacity and thus have limited reserve capacity for additional power flows during emergency

situations. Bottlenecks represent critical choke points in the transfer of power between interconnected networks.

The Network

Figure 4-2 shows an expanded view of Network 1. The distribution system forms a network of generators, substations, and major transmission lines. The network has five major generating stations (G1 through G5) that are responsible for generating power and four major transmission substations (S1 through S4) that distribute power. Generation, transmission, and power flows in each network are typically coordinated through a centralized operations and control facility. Network control centers have the primary responsibility of scheduling power purchases from generating units, allocating generation and loads to available transmission lines, ensuring network stability and reliability, and responding to emergencies.

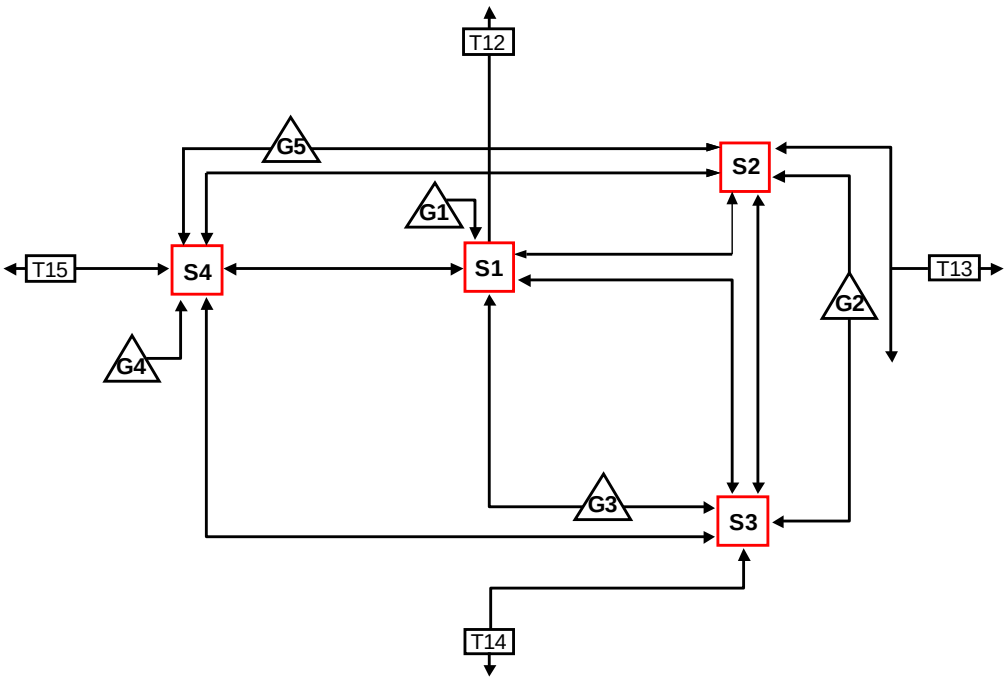

Figure 4-2. Generating Stations and Substations in Network 1

Elements of the Network and Region

This analysis focuses on four elements of the electrical grid: substations, transmission lines, supervisory control and data acquisition (SCADA) systems, and energy management systems (EMSs). Each represents a potential point of vulnerability and, therefore must be defined.

Substations

Substations (S1, S2, S3, and S4) are the transfer points for energy flows within the distribution grid. Each substation contains transmission line termination points, as well as circuit breakers and bus bars that interconnect the transmission lines with various circuits. Major substations contain transformers that reduce intertie transmission line voltages to network transmission levels. Each substation contains metering equipment, protection relays, and

switching circuits that control the operation of the connected generation, transmission, and distribution supplies.

In Figure 4-2, substation S1 contains: monitoring, control, and protection circuits for all power output from generating station G1; part of the power output from generating station G3; network transmission line connections to substations S2, S3, and S4; and regional transmission line interconnection T12.

Transmission Lines

Transmission lines are conduits that transfer energy throughout the grid. Because of their importance to system operation, this assessment focuses primarily on EHV transmission lines that transmit energy from individual generators to the major substations in each network.

Substation S1 contains the following six transmission line connections: line G1-S1 connects the output from generating station G1; line G3-S1 connects the output from generating station G3; lines S1-S2, S1-S3, and S1-S4 connect to the other substations in the network; and line T12 is the regional intertie to Network 2.

In practice, each transmission line typically contains two or more parallel circuits, either mounted on overhead towers or routed underground. Because land space available for EHV transmission corridors is often limited, several transmission lines may be routed through the same right-of-way. For example, transmission line T12 is the long distance tie line to Network 2. However, lines G1 and T12 are located in a common right-of-way for part of their route to substation S1. Similarly, lines S1-S2 and S1-S3 leave substation S1 together before they split.

Supervisory Control and Data Acquisition Systems

Each network SCADA system provides integrated parameter monitoring, data processing, and automatic control of circuit switching, load smoothing, and regulation of voltage and frequency throughout the network. The SCADA system also provides status displays for all major equipment and transmission lines, parametric trends, alarms, and a manual control interface for the load-control center operators.

The SCADA "oversees" the network and responds to changing conditions. For example, if generating station G1 trips off line, the consequential voltage and frequency fluctuations may require rapid, active circuit switching to route additional power to substation S1. The SCADA system automatically controls energy transfers by using appropriate circuit breakers and increases output from the remaining generators to compensate for lost generating capacity. If fluctuations cannot be stabilized, the SCADA system implements preprogrammed automatic protection protocols to separate the connections to substations and restore stable conditions throughout the remainder of the network. Similar supervisory and control functions are also performed by SCADA systems at the regional level.

Energy Management Systems

An EMS can be loosely thought of as providing input to the SCADA control system. EMS determines the most cost-effective configuration of power production, transmission, and distribution throughout the network, considering the required criteria for system stability, safety,

and reliability. An EMS typically provides the fundamental information and computation capability to perform real-time network analyses, to provide strategies for controlling system energy flows, and to determine the most economical mix of power generation, power purchases, and sales.

For example, if generating station G4 trips off line, EMS will determine if it is more cost-effective to increase output from generating station G5, to start local peaking units (auxiliary units to supplement high network demands), to increase energy flow through interconnection T15, or to implement other options. The strategy depends on the system status at the time of the transient and preprogrammed protocols for rapid recovery of stable load flows at the lowest available cost for emergency replacement power.

Particulars of the Example

For illustrative purposes, several particular system characteristics are defined:

1. The generating capacity in Network 1 is not sufficient to meet load demands during periods of peak energy usage (e.g., summer weekdays).
2. Most customers in Network 1 are supplied through connections to substations S1, S2, and S3.
3. Substation S4 serves largely as an EHV transmission intertie and carries only a small fraction of the total network distribution load.
4. Interconnection T12, the primary intertie between Network 1 and the region, is a potential bottleneck.

Step 2. Identify and characterize the "sources of danger," that is, the hazards (e.g., stored energy, toxic substances, hazardous materials, acts of nature, sabotage, terrorism, equipment failure, combinations of each, etc.).

CHARACTERIZING THREATS

Once a system is defined, the hazards associated with it can be identified and characterized. In the risk sciences, the word "hazard" is usually defined as "a potential source of danger or damage" but does not necessarily imply the infliction of damage. A risk scenario is a sequence of events that links the hazard to the final damage state. For example, a chemical plant with an inventory of toxic chemicals can contain a variety of hazards; but only through risk scenarios (i.e., accidents or malicious acts) can the hazards be linked to or manifested as an actual damage state.

For this example, the source of danger is defined as a potential terrorist action. Specific scenarios that will be developed further in the example are (1) a physical attack on the electrical grid; and (2) a complementary simultaneous cyberattack on the electrical grid.

Step 3. Develop terrorist attack scenarios to establish levels of damage and consequences.

CONSTRUCTING SCENARIOS

Scenario development, the fundamental building block of every risk assessment, follows a structured format that answers two of the triplet questions: what can go wrong?, and what the consequences would be? A variety of logic and analytical tools are used to develop scenarios. These include event-sequence diagrams (ESDs) that display important elements of the evolving scenario, failure modes and effects analyses (FMEAs) that tabulate possible contributing causes, and event trees or fault trees that display functional and logical relationships among threats, targets, vulnerabilities, and consequences.

Two common methods are used for scenario development: one involves going forward from an initial disturbance of the system; the other works backward from the undesirable end-state:

1. Given a set of initiating events, the structuring of scenarios is done so the end-state (the damage state or undesired event) of each scenario is the condition that terminates the scenario. This approach is used for full-scope risk assessments that trace a system upset from initiation to final impact on the system. Scenarios constructed in this way form what is called an event tree.
2. Given an end-state (the undesired event), project backwards to determine the potential scenarios that could cause the end-state. This approach yields what is called a fault tree.

These methods can be used together to construct an encompassing set of risk scenarios. Obviously, a comprehensive examination of the electrical grid vulnerabilities might identify a great number of possible threat scenarios for a particular set of consequences or damage levels. It is impractical in this example to demonstrate a complete risk assessment of all possible damage conditions. Therefore, we define a small number of possible scenarios and link them to defined damage levels.

In Step 2, the source of danger (the terrorist action) was defined and potential threats were identified. For this step, six potential end-states are defined and linked to initiating events through the scenario development process:

- Damage Level 0 (no damage) – no significant network or regional power outages
- Damage Level 1 – transient outage to Network 1
- Damage Level 2 – transient outage to the region (and Network 1)
- Damage Level 3 – long-term outage to Network 1
- Damage Level 4 – long-term outage to Network 1 and transient outage to the region

- Damage Level 5 – long-term outage to the region (and Network 1)

Damage Level 0 (included for analysis completeness) accounts for the possibility that the terrorist may fail to cause any significant damage. Actions that prevent or effectively mitigate an attack scenario result in Damage Level 0.

For the purpose of this example application and for simplicity, transient damage means a complete loss of power for a period of 4 hours to 24 hours. Long-term damage means a complete loss of power for more than 24 hours. For example, Damage Level 1 means that Network 1 (and only Network 1) experiences a power outage of up to 24 hours. Damage Level 4 means that Network 1 experiences a power outage of more than 24 hours and the entire region experiences an outage of up to 24 hours. The damage levels are used primarily to focus the scenario construction process and to show that a clear definition of the undesired consequence is critical to a structured risk assessment.

In this example, the damage levels do not explicitly account for health and safety consequences, but a sustained outage of electric power would clearly cause chaos and helplessness, especially in an urban environment. Depending on the duration of the outage and the interdependencies of infrastructures, the consequences could be catastrophic. Cascading events could lead to the loss of: transportation systems, clean water, sanitation, health care, security, and food supplies. Based on the damage conditions considered in the example, it would be possible to assess health and safety consequences for a specific urban setting using the same techniques.

Structuring the Scenarios

The scenarios show how specific damage levels can result from physical attacks on the system hardware, cyberattacks on system controls, and combinations of these attacks. First, a potential physical attack is discussed to illustrate how an event is generated. Second, a cyberattack is presented. The cyberattack may either initiate additional failures or further compound the effects of the physical damage.

Physical Attack on Network 1

Numerous physical methods could be used to damage equipment at each substation with varying degrees of damage to the network and the region. For example, carbon fibers, Mylar strips, or other contaminants could be sprayed over buses and transformers to cause severe short circuits. Explosives could be used to destroy key transformers, circuit breakers, and bus sections. Attackers could also damage circuit breaker controls at substation operating panels.

To help model this scenario, we will assume that a threat assessment uncovered a high likelihood that detailed information about the electrical grid has been made available to terrorists. To illustrate the method, substation S1 is analyzed first because it controls the full output from generating station G1, part of the output from generating station G3, and, most important, the termination of the key regional interconnection T12.

To generate the attack scenarios, five sequential questions are asked:

1. Does the attack succeed? Success means that substation S1 is physically attacked and disabled.
2. Do all of the other generating stations in Network 1 fail? Electrical grids are typically designed so that one substation can trip offline without destabilizing the entire network. However, it is conceivable that a fault-initiated clearance of all circuits at substation S1 could cause a sufficient drop in voltage and frequency to initiate automatic load shedding and circuit isolation at all other substations, thereby causing all remaining generators to trip off line. Therefore, our model must account for this possibility.
3. How long does the Network 1 outage last? This will depend on how quickly contingency plans can be implemented to enable the network to recover.
4. Does the transient propagate through the region? In the event of a physical attack that destabilizes Network 1, it is very likely that the regional protection signals would automatically open the remaining interconnections (T13, T14, and T15) to prevent the transient from propagating to adjacent networks. However, the possibility of failures that cause cascading damage at the regional level must be considered for a comprehensive analysis of transient and long-term outages.
5. How long does the regional outage last? This will depend on how quickly contingency plans can be implemented to enable the system to recover. It is assumed that if there is a transient network outage, then the maximum time for a regional outage is also transient.

Figure 4-3 shows the systematic thought process used to develop the attack scenarios and to assign their consequences to the damage levels. Branches may be added to account for other protective barriers in each system. The purpose of this exercise is to create a comprehensive framework for identifying vulnerabilities and in turn make better decisions. Figure 4-3 illustrates how the scenario development process can be used to represent a complex scenario. A full-scale risk assessment would detail the effects from attacks on each substation, as well as multiple substations at once.

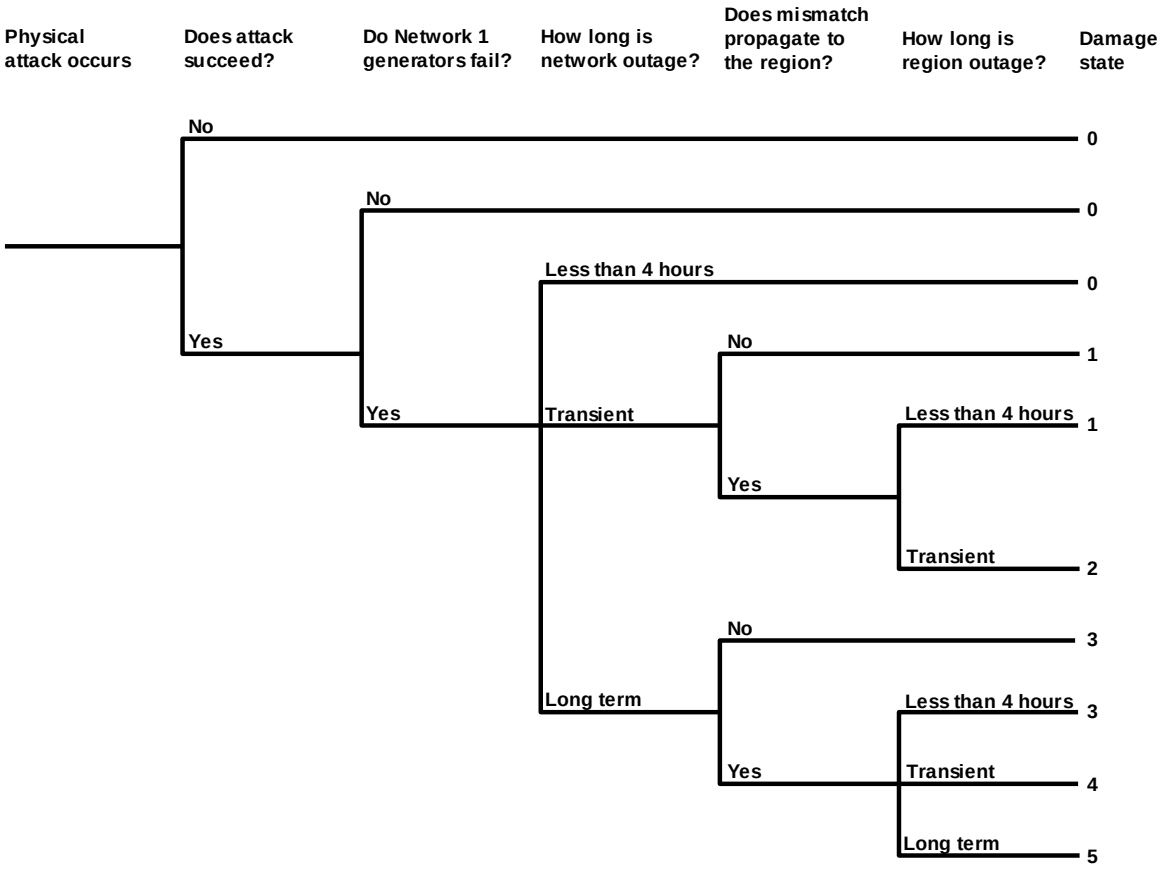

Figure 4-3. Thought Process for Attack Scenarios

Cyberattack in Conjunction with a Physical Attack

Given that a physical attack has destabilized Network 1, how might a terrorist prevent isolation of the network and allow the transient to propagate into the regional grid? One method would be to coordinate the physical attack with a cyberattack that keeps the circuit breakers closed at the network-region interties.

The growing complexity of the electric power grid, coupled with economic incentives for trading energy across regions, has significantly increased reliance on computerized control systems and data communication networks to control the components of the electrical grid. This leads to a potential vulnerability that can be exploited by terrorists. A cyberattack may be attractive to terrorists for many reasons:

- A cyberattack on the electric power grid would not require a physical presence in the United States. The attack could be planned, coordinated, and carried out from almost anywhere in the world where there is a connection to the Internet, thus eliminating the security risks and expenses of infiltrating human agents into the United States where they and their plans might be discovered.

- Significant damage could be done with minimal investment. This same logic has been hypothesized as a motivating factor behind the September 11 attacks—the attacks were planned to produce the most damage with the least investment (e.g., a high return on investment for terrorist monetary and personnel resources).
- The terrorists would use our own resources to attack us. This is also compatible with one of the hypothesized characteristics of the current threat—the major resources for an attack are supplied by the target nation. The open, unregulated nature of the Internet in the United States provides a wide-open pathway to targets. This Internet pathway provides not only a way to reach SCADA systems, but is also an invaluable resource for identifying potential targets and providing technical information critical to the success of an attack.

As Figure 4-3 shows, power outages in Network 1 may propagate into the regional grid if the regional SCADA emergency protection and control functions are disabled. Thus, one possible way for a terrorist to cause a regional outage is to ensure that two successive events occur: (1) a very large power mismatch in Network 1 must be created (e.g., by a physical attack); and (2) the initiating transient created in Network 1 must propagate through and disable the region (e.g., by a cyberattack).

One possible way to achieve Damage Level 4 is to ensure that the power mismatch created by a physical attack cannot be quickly corrected by combinations of available generation and automatic load shedding in Network 1 or by automatic supplies from the interconnected regional grid. After Network 1 is brought down, additional steps would be necessary to ensure that the Network 1 failures cascade throughout the regional grid. Thus, intruders must override or block the regional SCADA protection and control systems that contain the frequency stabilization, load shedding, and islanding protocols. If the major regional interties remain connected to the faults in Network 1, the entire grid will quickly collapse. Individual network protection and control systems will attempt to maintain stable power flows within each of the other networks. However, if a network depends heavily on bulk power flows from the regional grid, it is very likely that the internal network control systems will not stabilize voltage or frequency. Widespread automatic shedding of loads and generation will then cause additional outages and exacerbate instabilities in other networks along the line. Of course, causing this level of regional damage would typically require more resources and coordination than an attack that affects only Network 1.

Step 4. Adopt risk metrics that reflect the likelihoods of different attack scenarios in terms of target and collateral damage and quantify the scenarios based on the totality of relevant evidence.

RISK ASSESSMENT

The triplet definition of risk is the framework for measuring risk. Risk is measured in terms of scenarios (what will happen), likelihood (how likely it is to happen), and consequences (what the results would be). Risk is not a number, but a collection of numbers, or more precisely

a collection of curves that display scenarios, likelihoods, and consequences. The so-called "risk parameter" is usually expressed as the frequency with which an undesired event occurs. Since this frequency is never known exactly, our state of knowledge about the numerical value of this frequency is expressed as a probability curve against the possible numerical values of the frequency. This probability curve is used in the Bayesian sense and expresses our state of knowledge about the frequency, based on all the relevant evidence available. Probability interpreted in this way embodies the notion of uncertainty. The undesired event(s) can be a fixed level of damage, such as the total destruction of a building, or a varying parameter, such as the number of fatalities or injuries with probability as a parameter. Dollars are also a widely used parameter for measuring risk. In many situations, combinations of risk measures are used.

In this section we will examine how the model for assessing the risk of power failures at the network and regional levels is constructed and, separately, how each type of attack is modeled and quantified. The quantification of the attack scenarios follows the process described in Chapter 3. The first step of the process is to develop a model that provides a framework for systematically evaluating the causes, frequencies, and consequences of each undesired condition. Experience has shown that the top-down perspective (employed here) is the best way to ensure that the analyses are complete. The scope of the model must be broad enough to account for all possible causes and all possible consequences. The model must also be sufficiently detailed to support realistic engineering evaluations of various threats and vulnerabilities and to provide clear information about the contributors to each undesired event. The model must support quantitative analysis of each potential contributor, including rigorous treatment of uncertainties throughout the analysis process.

The parameter selected for measuring risk is based on the success rate of different levels of damage. The probability of frequency concept introduced in Chapter 3 is a convenient parameter for calculating risk because it not only represents the frequency with which a specific consequence may occur, but it also communicates the analyst's uncertainty in that frequency and, therefore, in the risk. In this example, the success rate for different levels of damage (a form of frequency) was chosen as a convenient parameter. Thus, the probability of the success rate for achieving different consequences, or damage levels, is the basis for measuring risk.

Top-Level Event Tree

Figure 4-4 shows a simplified top-level event tree that may be used to quantify the levels of damage in this example. The following items briefly summarize the scope and definition of each top event listed in the figure.

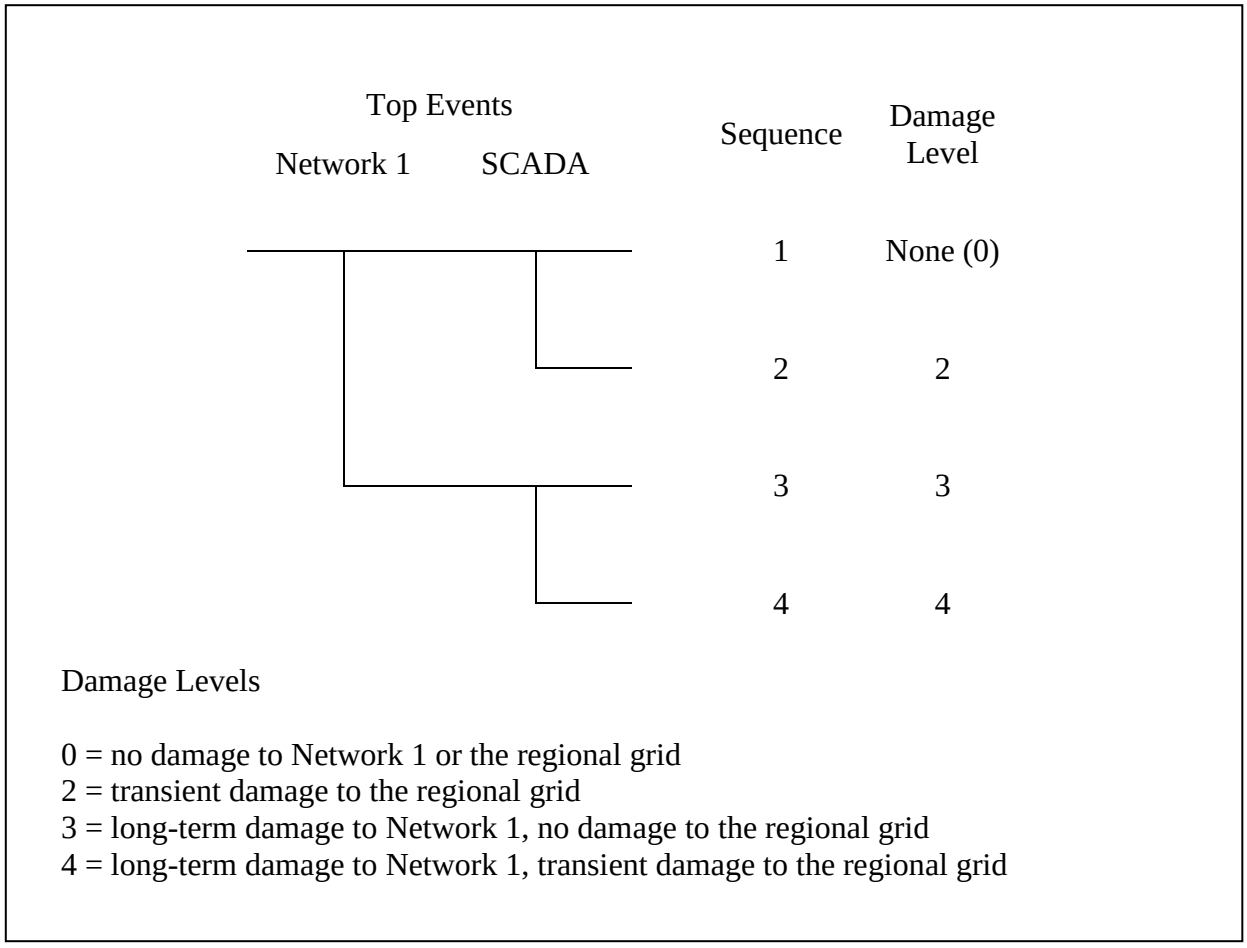

Figure 4-4. Top Level Event Tree for Grid Damage

Network 1

The Network 1 top event represents the success rate for attackers damaging sufficient equipment in Network 1 to cause a long-term power outage. The horizontal path from the Network 1 top event occurs if the attackers do not disable enough equipment to cause a network power outage. The failure path from the Network 1 top event (the vertical path in the event tree) occurs if the attack results in long-term damage to network power supplies.

SCADA

The SCADA top event represents the success rate for intruders initiating a cyberattack that causes short-term power failures throughout the regional grid. The horizontal path from the SCADA top event occurs if the intruders do not disable the regional grid. The failure path from the SCADA top event (the vertical path in the event tree) occurs if the intruders cause a regional power outage.

Possible Outcomes

Sequence 1 occurs if the attackers do not achieve any of their objectives. Even if there are some localized power outages in Network 1 or in portions of the regional grid, the outages are not severe enough or of long enough duration to satisfy the damage criteria of concern for the

analysis. Sequence 1 terminates in a condition considered to be a functional success of the regional and network power supplies; it is assigned to Damage Level 0.

Sequence 2 occurs if intruders successfully initiate a cyberattack on the regional SCADA systems causing them to send out anomalous protection and control signals causing widespread, short-term power outages throughout the grid, including outages in Network 1. However, the local attackers are not able to cause sufficient damage to equipment to prolong the outages in Network 1. Sequence 2 terminates in a condition equivalent to Damage Level 2.

Sequence 3 occurs if the attackers cause sufficient damage to Network 1 to cause widespread, long-term power outages throughout a large portion of the network, but no disruption in regional power supplies. Sequence 3 terminates in a condition equivalent to Damage Level 3.

Sequence 4 occurs if the attackers achieve all of their objectives. The local attackers cause sufficient damage in Network 1 to cause widespread, long-term power outages throughout a large portion of the network. A successful cyberattack on the regional control systems also prevents the normal operation of protection signals or causes other active signals that disrupt regional power supplies. Sequence 4 terminates in a condition equivalent to Damage Level 4.

The top-level event tree in Figure 4-4 is logically complete and provides a framework for evaluating the success rate of each potential level of damage. In practice, however, it is often necessary to increase the level of detail in the supporting analyses to examine the threats, vulnerabilities, and causes that may contribute to each undesired condition. The increased detail facilitates a more systematic evaluation of each potential cause of failure and provides a logical framework for assessing the effectiveness of specific mitigation measures. The detailed evaluations also often reduce the uncertainties inherent in approximate, high-level estimates or identify the most important sources of uncertainty in each estimate.

Modeling the Physical Attack: Top Events

The event tree in Figure 4-5 is a more detailed analysis of the Network 1 top event. The expanded logic includes more details about attacks on the three critical substations in Network 1 and the corresponding likelihoods of a long-term network power outage. The scope and definition of each top event listed in the figure are summarized below.

SUB S1

The SUB S1 top event represents the success rate for the attackers destroying sufficient equipment in substation S1 to disable its power generation and transmission interconnections. The horizontal path from the SUB S1 top event occurs if the attackers do not achieve their goal; that is, the substation may be partially damaged, or the effects may temporarily disrupt power. However, the damage is not sufficient to incapacitate the major interconnections for more than 24 hours. The failure path from the SUB S1 top event (the vertical path in the event tree) occurs if the attackers cause enough damage to equipment to disable substation S1 for an extended period of time.

SUB S2

The SUB S2 top event is similar to the SUB S1 top event. It represents the success rate for the attackers destroying enough equipment in substation S2 to disable its power generation and transmission interconnections.

SUB S3

The SUB S3 top event is similar to the SUB S1 top event. It represents the success rate for attackers destroying enough equipment in substation S3 to disable its power generation and transmission interconnections.

NET

The NET top event represents the conditional success rate for each level of substation damage causing an extended power outage throughout Network 1. This success rate depends on the specific combination of substations that are damaged, their generation and transmission interconnections, and the network loading conditions at the time of the attack. The success rate for a consequential failure of the NET top event is different for each combination of damage conditions to substations.

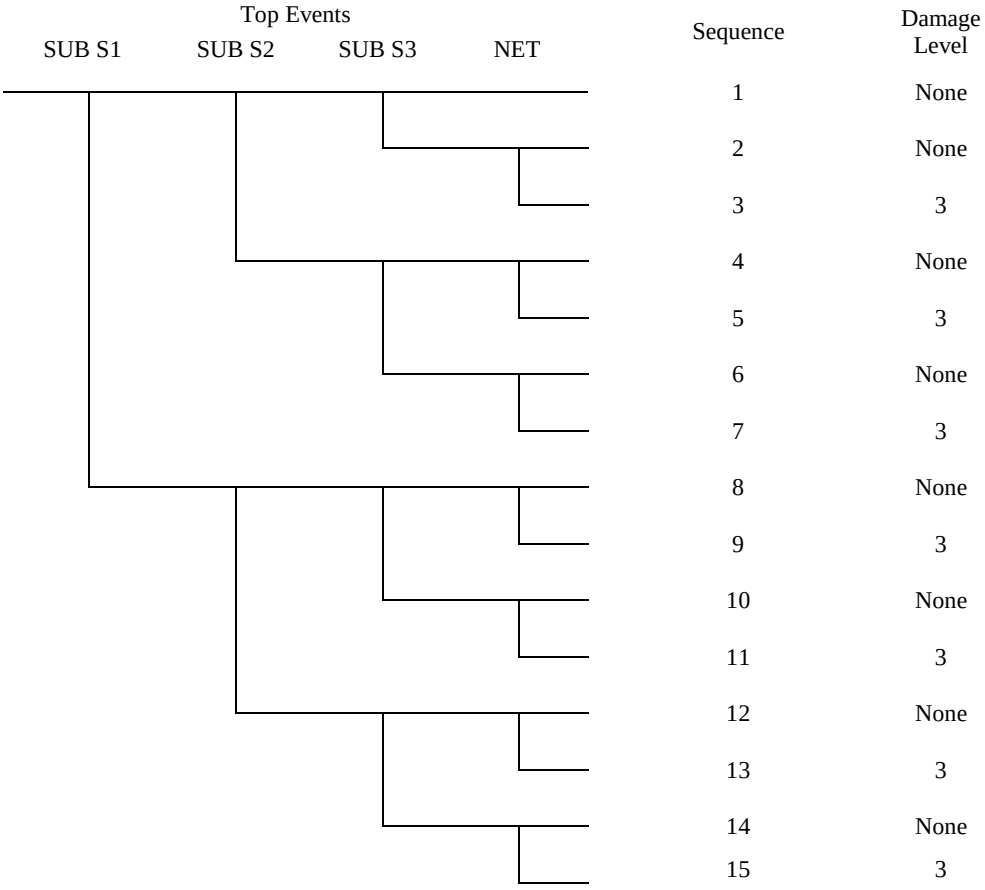

Figure 4-5. Event Tree with Increased Detail for the Network 1 Event

Modeling the Physical Attack: Possible Outcomes

Sequence 1 in the event tree occurs if the attackers do not cause enough damage to incapacitate any of the three critical substations. Short-term, localized power disruptions may occur in some areas, but the outages are not of sufficient severity or duration to satisfy the damage criteria of concern for the analysis. The success path from the NET top event also occurs if the attacks do not inflict enough damage to cause widespread extended outages throughout the network. Thus, sequences 1, 2, 4, 6, 8, 10, 12, and 14 end in a condition considered to be functional success of the network power supplies.

Sequence 3 in the event tree occurs, if the damage to substation S3 is severe enough to cause prolonged power outages throughout a large portion of Network 1. The failure path from the NET top event occurs whenever the achieved level of substation damage is severe enough to cause widespread extended outages throughout the network. This condition occurs in sequences 3, 5, 7, 9, 11, 13, and 15 and is equivalent to Damage Level 3.

In practice, for a more detailed evaluation of the possible contributions to long-term network outages, the Network 1 top event in Figure 4-4 can be replaced by the entire event tree in Figure 4-5. Of course, other types of logic models can be used to accomplish the same goal (e.g., a fault tree that is logically equivalent to Figure 4-5). More detailed models may be developed to further subdivide and evaluate the various threats and vulnerabilities that contribute to each top event. For example, numerous potential attack scenarios with specific requirements for attacker resources and corresponding likelihoods of success may be examined for substation S1. Coordination strategies for attacks on multiple targets may also be examined, which may introduce important dependencies among the analyses for each substation. For the purposes of this example, the level of detail is shown only for the integration of Figures 4-4 and 4-5.

Evaluation of Threats and Vulnerabilities in Terms of the Supporting Evidence[4]

The most important function of a risk model is to organize the problem logically and provide a structured format for the systematic examination and evaluation of contributing threats and vulnerabilities. Figures 4-4 and 4-5 provide a logical framework with enough detail to perform a top-level evaluation of the risk associated with each level of damage considered in this example. The most difficult part of the risk assessment process is the development of realistic, quantitative estimates for the likelihood of each potential failure, including consistent evaluations of the uncertainties in each estimate.

Quantifying the Physical Attack

The following section summarizes quantitative estimates developed specifically for the physical attack on Network 1. Further on, quantitative estimates will be developed for the cyberattack at the regional level. Although these estimates are necessarily simplified and are not derived from detailed analyses of any particular electrical network or regional control system, they illustrate the types of analyses, thought processes, and inputs that are typically developed to support the risk assessment process. In its simplest form, as in this example, inputs may be

[4] To simplify the process for this example the data consist of estimates by experts rather than data from searches and the application of Bayes Theorem for inferring probability distributions.

based on the experience and judgment of experts. Even though these high-level screening analyses are typically only approximate and often include large uncertainties, they are useful for focusing attention quickly on specific elements of the problem or parts of the analysis that merit more careful, more detailed evaluation. Additional detail may then be added to the models for those elements, and their supporting analyses refined to identify the most important contributing causes or to reduce the initial uncertainties.

Top events SUB S1, SUB S2, and SUB S3 in Figure 4-5 represent the likelihood that attackers would destroy enough equipment in each substation to disable its generating supplies and transmission interconnections. In this model, each critical substation is assigned a different vulnerability to attack.

Substation S1. It is assumed that substation S1 is located in an urban environment and is the most heavily protected of the three substations. It may be surrounded by protective walls, may be continually manned by utility personnel, and may be checked by local police during their normal neighborhood surveillance patrols.

Substation S2. It is assumed that substation S2 is located in a suburban or partially rural environment and is the least protected of the three substations. It may be surrounded by a chain link fence, may not be manned, and may not be subject to routine surveillance by local police.

Substation S3. It is assumed that substation S3 is located in an urban environment but is only partially protected. For example, it may be surrounded by protective walls and checked by local police during their normal neighborhood surveillance patrols, but it may not be continually manned.

A simple probability distribution can be developed to assess the likelihood that attackers could successfully enter each substation and cause extensive damage to critical transformers, circuit breakers, buses, and controls. The histogram in Figure 4-6 applies to substation S1.

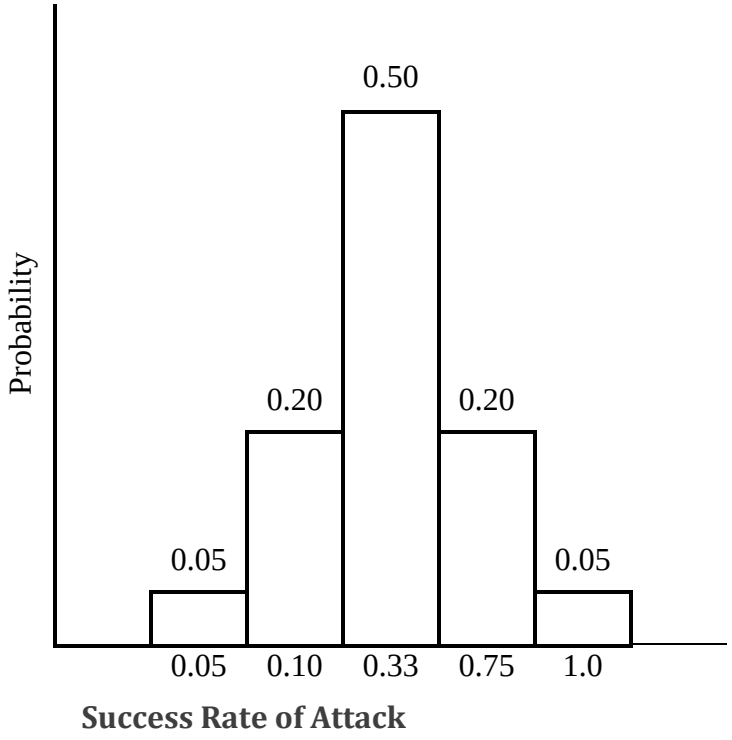

Figure 4-6. Histogram Showing Success Rate of an Attack on Substation S1

This simple histogram does not rigorously display the discrete probability in boundaries over the full range of the cumulative probability distribution function. Nevertheless, it is useful for demonstrating the fundamental concepts that would be used for a more numerically rigorous representation of the uncertainties. The sample histogram shows the following information:

- There is a 5 percent probability that the attackers would succeed in 5 percent of their attacks on substation S1 (i.e., that 1 of 20 attacks would be successful).
- There is a 20 percent probability that the attackers would succeed in 10 percent of their attacks on substation S1 (i.e., that 1 of 10 attacks would be successful).
- There is a 50 percent probability that the attackers would succeed in 33 percent of their attacks on substation S1 (i.e., that 1 of 3 attacks would be successful).
- There is a 20 percent probability that the attackers would succeed in 75 percent of their attacks on substation S1 (i.e., that 3 of 4 attacks would be successful).
- There is a 5 percent probability that the attackers would always succeed in their attacks on substation S1 (i.e., that every attack would be successful).

According to these estimates, the mean likelihood of a successful attack on substation S1 is approximately 0.39 (i.e., approximately 10 of 26 attacks would be successful). These estimates are obviously not derived from detailed models of specific attack scenarios or from a

detailed evaluation of the specific substation vulnerability to each attack. However, these types of estimates can be developed relatively easily, based on information from experts familiar with potential attack strategies, resources, and specific vulnerabilities of the target. If these preliminary results showed that attacks on substation S1 were potentially important to one of the undesired damage levels, more extensive analyses would be justified.

Table 4-1 summarizes the estimates of a successful attack on each substation, considering its specific vulnerabilities. These estimates account for the conditional likelihood of success after an attack is launched, but they do not explicitly account for pre-attack planning to identify key targets, evaluate critical network-loading conditions, develop logistics for the attack teams, etc. These factors would obviously also influence the overall likelihood of a successful attack, especially a coordinated offensive on multiple targets. In a more detailed analysis, these factors could be included as additional inputs to the models for each substation, or they could be evaluated in a separate part of the risk model that specifically examines the planning, resources, and logistics of the attack.

TABLE 4-1. Estimated Success Rate of an Attack

Substation	Probability					Mean
	0.05	0.20	0.50	0.20	0.05	
S1	0.05	0.10	0.33	0.75	1.0	0.39
S2	0.75	0.85	0.90	0.95	1.0	0.90
S3	0.10	0.30	0.50	0.80	1.0	0.53

Evaluating the NET Top Event

An evaluation of the NET top event in Figure 4-5 accounts for the conditional likelihood that each level of substation damage would cause an extended power outage throughout Network 1. The description of the network indicates that substation S1 is the most important one because it controls the full output from generating station G1, part of the output from generating station G3, and the termination of key regional interconnection T12. Substation S2 is next in importance because it contains the connections from generating stations G2 and G5, which are not directly connected to substation S1. Substation S3 is the least important of the three critical substations.

The network is designed to withstand the complete loss of any one substation under normal loading conditions. However, under severe loading conditions, attack-initiated faults might cascade to other substations and generating units. Therefore, the models for the NET top event must assign a likelihood of network failure after any combination of substations is damaged. In this simplified example, these conditional likelihoods are expressed by the probability histograms summarized in Table 4-2. Of course, in a more detailed analysis, additional supporting information for these estimates could be derived from dynamic load-flow simulations, models of system response, interviews with network operations personnel, etc.

TABLE 4-2. Conditional Success Rate for a Network 1 Failure

Damaged Substations	Probability					Mean
	0.05	0.20	0.50	0.20	0.05	
S1	0.05	0.10	0.25	0.50	0.75	0.29
S2	0.01	0.10	0.15	0.25	0.50	0.17
S3	0.01	0.05	0.10	0.20	0.25	0.11
S1 and S2	0.25	0.50	0.75	0.90	1.0	0.72
S1 and S3	0.10	0.25	0.50	0.75	1.0	0.51
S2 and S3	0.05	0.10	0.25	0.50	1.0	0.30
S1, S2, and S3	0.90	0.92	0.95	0.98	1.0	0.95

Modeling the Cyberattack: Intrusion into Regional SCADA Control

The cyberattack scenario outlined here takes place over a three-week period. Although it is possible the attack could be orchestrated in much less time, it is assumed that the characteristics of a September 11 type attacks (e.g., cautious and careful planning) would be in operation; thus a three-week timeline might be more typical. The cyberattack is divided into five phases: (1) discovery; (2) launch platform acquisition; (3) target selection; (4) target reconnaissance and compromise; and (5) initiation of an actual attack on the electric power grid.

Discovery Phase

The discovery phase of the operation begins with the identification of potential targets and the assembly of critical information about them. Actors with very little computer knowledge could carry out this phase of the attack, and there is a good chance that the activities during this phase would be carried out by individuals other than those who would be responsible for the final attack. This would compartmentalize resources and protect higher level technical operatives from possible exposure and loss.

The first step would be to identify potential targets via the Internet. This could be done using one of hundreds of search engines by typing in keywords, such as “power company,” “electric power,” “power and light,” or other common phrases associated with electric utilities. In just a few hours, a large number of U.S. electric utility companies could be identified. Alternatively, the names of every private and municipal electric utility in the United States could be collected in electronic format in less than five minutes from a publicly available government website.

The next step in the discovery phase would be to find the computer systems of the electric utility companies that are connected to the Internet. Like most institutional entities with a presence on the Internet, electric utility companies have registered and reserved large ranges of IP addresses. Registered IP addresses are unique to the registered entity; they are the “electronic address” by which they can be reached from anywhere else on the Internet.

One efficient way to collect these addresses would be to access one of thousands of "whois" engines on the Internet. In just seconds, these publicly available search engines can search millions of IP address registration records and identify the addresses associated with keywords, such as "XYZ Power and Light" or other specific electric utility company names collected in the first part of the discovery process. The IP addresses that surface from these "whois" searches could then be cut and pasted into a local document, such as an Excel spreadsheet on the discovery team's computer. Once this has been accomplished, tens of millions of unrelated Internet addresses would have been eliminated, and a database of potential electric utility computer systems would have been assembled. It is likely at this point that the discovery team would encrypt their electric utility system database, burn it on to a compact disc, and hand it off to a courier who would physically carry it to the attack team. This would prevent it from being intercepted by the National Security Agency or another intelligence-gathering organization.

Launch Platform Acquisition

For security reasons, the actual attack team would most likely be located in a country other than the one in which the discovery team resides. The attack team would probably include several intermediate-level computer users and one expert computer hacker. Their first task would be to compromise a series of computers from which to launch the attack. Computer attacks are typically carried out through a series of computers, which makes it very difficult to trace the source of the attack, if it is even discovered. The attack team would prowl computer networks in countries where computer security is poor or nonexistent. Using autorooters, port scanners, and other tools that are readily available on the Internet, they would scan computer networks in these vulnerable countries looking for computer systems with vulnerabilities that could be exploited. Once found, the computers would be compromised; the attackers would arrange administrative privileges on these machines and then go dormant, covering their tracks by deleting log entries and using other stealth techniques. In this manner, they would build a set of computer systems from which they could launch their cyberattacks remotely.

Target Selection

The actual portion of the electrical grid selected as a target might depend on an a priori selection of targets by higher-level operatives in the terrorist organization to coordinate with a physical attack on the power grid or even on another interdependent infrastructure target. However, the terrorist organization might also settle for a target of convenience and leave the decision up to the attack team.

In any event, once an electric utility had been selected as a target, the attack group would activate some of the computers exploited in the platform acquisition phase, transferring the autorooter and port scanning tools to the compromised computers. Next, those tools would be used against the utility's range of IP addresses in the discovery team database. Many autorooters are sophisticated and automated—that is, they can try multiple attack strategies against a large range of machines. When they are successful, they can install a number of surveillance/reconnaissance tools that would automatically cover up any sign that the utility computer had been compromised.

Several classes of commercially available products are designed to protect against these kinds of attacks. A number of computer firewall products are designed to recognize and deflect attacks like the one described above by restricting all incoming and outgoing network traffic unless the administrator of the firewall designates it. A second class of security devices, intrusion-detection systems, monitors incoming and outgoing network traffic for digital signatures of known cyberattack tools and ploys. Although these security techniques are often effective, they are not 100 percent effective; in fact they are often compromised by mis-configurations by the administrator. An acute shortage of well-trained computer security professionals is a contributing factor to the problem of computer security.

Target Reconnaissance/Compromise

The initial electric utility computer system that was compromised in the previous stage would most likely be an administrative server, web server, or other computer not directly involved in the SCADA system, and, therefore, not the final target of the cyberattack. Cyberattacks with preplanned goals or objectives, such as the one in our terrorist scenario, usually use "attacks by increment" strategies.

In this phase, the computer system compromised in the previous stage would be used as the home base for the cyberattackers who would attempt to find out the purpose of the compromised computer and then assess the number of other computers in the network that "trust" the compromised computer and to what extent. They could then use these trust relationships to inspect other computer systems on the network, as well as to discover other local networks. The cyberattackers might also install packet sniffers to listen in on network traffic for packets destined for ports specific to a particular SCADA software system. Once they found SCADA port traffic, they could identify the computer systems being used as SCADA systems.

If the compromised computer does not provide a pathway to the SCADA network, the attackers would go back to the previous phase and attempt to compromise another externally visible computer system in the utility company's IP range. Another possible outcome might be that another vulnerable computer system (but not a SCADA controller) on a connected, but different network in the utility system would be identified; this computer would also be compromised, and reconnaissance could then be initiated from a newly compromised machine.

Initiation of Attack

The final step would involve compromising one or more of the computer systems that run the SCADA system. These systems would be attacked using the same autoroot and exploit tools that gained access to the initial computer in the electric utility. Once the SCADA system was compromised, the amount of damage inflicted on the components of the power grid reachable by the compromised SCADA system would depend on the attack team's knowledge of electric power systems.

Ideally, one member of the attack team would be a power engineer trained in the basics of power generation and distribution systems. The damage inflicted could be significantly increased by knowledge of the specific power system and components that would be under the control of the terrorist group.

TABLE 4-3 Stages of SCADA System Intrusion

Day	Event	Objective	Actors	Probability of Success	Probability Parameters	Choke Point	Probability of Successful Choke Point Intervention
1	Use internet search engine to find U.S. power companies.	Identify potential targets and power generation sites.	Low-level operatives.	Near 1.0	Presence of power company and generation site details on Internet.	No	Near 0
1	Search "whois" engine for names of power companies discovered above.	"Who is" records will contain IP address blocks assigned to the company, thereby drastically reducing the search space for exploit targets.	Low-level operatives.	0.8	Some power companies are listed but have blocked IP addresses; others have large blocks of IP addresses registered.	No	Small
3	Deploy autorooter exploit at foreign networks in Korea, India, etc. Identify vulnerable systems. Plant exploit to take control of N systems.	Create a network of exploited computers in countries with many poorly protected networks. Typical tactic is to gain root access to a number of machines and then connect to the target machine through multiple IP connections to hide the true IP address of the attacker.	Actors with modest to intermediate computer skills.	1.0	The number of vulnerable computers on the Internet, especially in certain parts of the world, for all purposes makes this an almost certain element of any attack.	No	
6	Deploy autorooter and portscan exploits on networks and computers captured on Day 3 against IP ranges of power company networks discovered on Day 1.	Look for vulnerable computer systems at U.S. power companies. These systems may be poorly protected web servers, administrative computers, or (if you're really lucky) a computer with direct SCADA duties.	Actors with intermediate to expert computer skills.	Near 1.0	Some vulnerable computers, especially in the administrative and web server classes, would be somewhere in the search space.	Unlikely	This should not be considered a choke point because there are too many entry points to ensure that all of them have been adequately protected.
8	Evaluate most likely targets from list of vulnerable power company computers found on Day 6. Pick top 3 or 4 targets in terms of attractiveness and deploy the appropriate exploit tool to gain covert control over the computer.	Gain covert control over the power company computer. The first objective is to determine what purpose the compromised machine serves to determine its potential as a launching pad toward the SCADA system. Exploit code may be encrypted or packet fragmentation may be used to avoid detection by firewall software.	Actors with intermediate to expert computer skills.	0.2	The actual probability of success depends on the security posture of the particular power company and the type of computer (web server, administrative machine, etc.) exploited, patches to firewall, and operating systems.	Yes	0.9 plus With the proper firewall/operating system software and system security, the probability of a breach can be kept to a minimum.

TABLE 4-3 Stages of SCADA System Intrusion

Day	Event	Objective	Actors	Probability of Success	Probability Parameters	Choke Point	Probability of Successful Choke Point Intervention
8	Execute routines to subvert logs that would tip off systems administrator of intrusion.	Intrusion detection avoidance either by a deployed intrusion detection system or a sharp-eyed computer systems administrator.	Actors with intermediate to expert computer skills.	0.4	Once a firewall has been defeated, it is likely the exploit used is one for which the intrusion detection system does not yet have a digital signature. Therefore, once past the firewall, an intrusion is more likely to go unnoticed.	Yes	0.9 plus A good intrusion detection system can make it difficult for terrorist groups without sophisticated computer knowledge to go undetected.
10	Examine the list of hosts, trusted hosts. "Sniff" packets of traffic going through the compromised machine.	Understand the role of the currently compromised computer in the power company's computer network. The next step is to explore the network the compromised computer is on to find other computers on the network as well as other networks to which it is connected. Note that this activity often occurs one or more days after the successful intrusion. Some initial research shows that attackers often lie low for a day or so after an intrusion and then return to see if their exploit is still present and viable (i.e., undiscovered).	Actors with expert computer skills to avoid detection during exploration of other networks on power company's computer infrastructure.	0.1	The probability here refers to the chance that the attacker will be detected during his exploration activities, which are risky because they involve multiple machines and may generate unusual traffic on other networks (e.g., IP addresses normally not seen passing traffic on a particular network).	No	

TABLE 4-3 Stages of SCADA System Intrusion

Day	Event	Objective	Actors	Probability of Success	Probability Parameters	Choke Point	Probability of Successful Choke Point Intervention
10 to 15	Explore the power company's computer network looking for evidence of SCADA activity.	Some of SCADA systems knowledge is required either directly or by following a "cookbook" set of SCADA indicators. These indicators would include looking for specific files in certain directories, certain processes that could be identified by doing something as simple as looking at the threads currently being executed on the machine, looking for traffic on certain ports, looking for particular hardware drivers for devices associated with SCADA hardware/software interfaces.	Actors with expert computer skills to avoid detection; direct or indirect expertise with power systems to identify markers that identify a SCADA system.	0.01	The opportunity to find a SCADA machine depends greatly on the level of security and connectivity within the company's computer networks.	Yes	0.9 plus Probably the best protection is air, that is, keeping the SCADA systems disconnected from other networks. Of course, the Internet will significantly reduce the odds of an attack.
16 to 17	Identify a SCADA computer and a careful process of attacking it with an exploit to gain covert control.	Gain control of a machine inside the SCADA system for intelligence gathering and for use as an exploit launching pad.	Actors with expert computer skills and at least some SCADA experience.	0.3	Once inside the SCADA network, it is likely that the attacker will find machines relatively less protected because of their "center perimeter" location, as well as because many of the characteristics of SCADA systems, such as mandatory fast response times to events, preclude extensive use of time-consuming encryption, packet inspection, or authentication.	Maybe	Intelligent security command and control agents that can isolate potential "bad guy" traffic with as little disruption to the power grid as possible may be able to make this a choke point.
17 to 21	Identify additional SCADA computers on the network and run exploits to gain control of them as well.	Gain control of as many SCADA control systems and devices as necessary to increase the amount of damage that could be done, as well as to reduce the probability that intervention by a power control systems operator would limit or prevent damage to the power grid.	Actors with expert computer skills and at least some SCADA experience.	0.7		Maybe	Intelligent security command and control agents that can isolate potential "bad guy" traffic with as little disruption to the power grid as possible could make this a chokepoint.

Quantifying the Cyberattack

The overall success rate for the SCADA top event in Figure 4-4 can be estimated from the evaluations of each step in the intrusion process and then combining the various event probabilities. The composite success rate of an intrusion that gains full control over the SCADA system is estimated to be approximately 1.3×10^{-5} per attempt (i.e., approximately 1 success in 75,000 attempts). This estimate is based on a probabilistic combination of the event probabilities of Table 4-3 and includes a very large uncertainty. For this example, the estimate by experts was used as the median value of a lognormal uncertainty distribution with an error factor of 10. This means that the experts were 90 percent confident that the likelihood of success would be within a factor of ±10 of the estimated value. The parameters of this uncertainty distribution are shown in Table 4-4.

Table 4-4. Probability Distribution of a Successful SCADA Intrusion (likelihood of success per attempted intrusion)

5^{th} Percentile	Median	95^{th} Percentile	Mean
1.3×10^{-6}	1.3×10^{-5}	1.3×10^{-4}	3.5×10^{-5}

Step 5. Assemble the scenarios according to damage levels, and cast the results into the appropriate risk curves and risk priorities.

Assembly

Once the individual scenarios have been quantified, they can be assembled into risk measures. This is a matter of combining all scenarios that terminate in a specific damage category. If the risk measure is a variable, such as fatalities, injuries, or dollars, then the process also involves arranging the scenarios in order of increasing damage and cumulating the probabilities from bottom to top.

A simple risk model that integrates the event trees in Figures 4-4 and 4-5 is quantified based on the supporting data summarized in the preceding section. The combined risk results for all damage levels are shown graphically in Figures 4-7, 4-8, and 4-9.

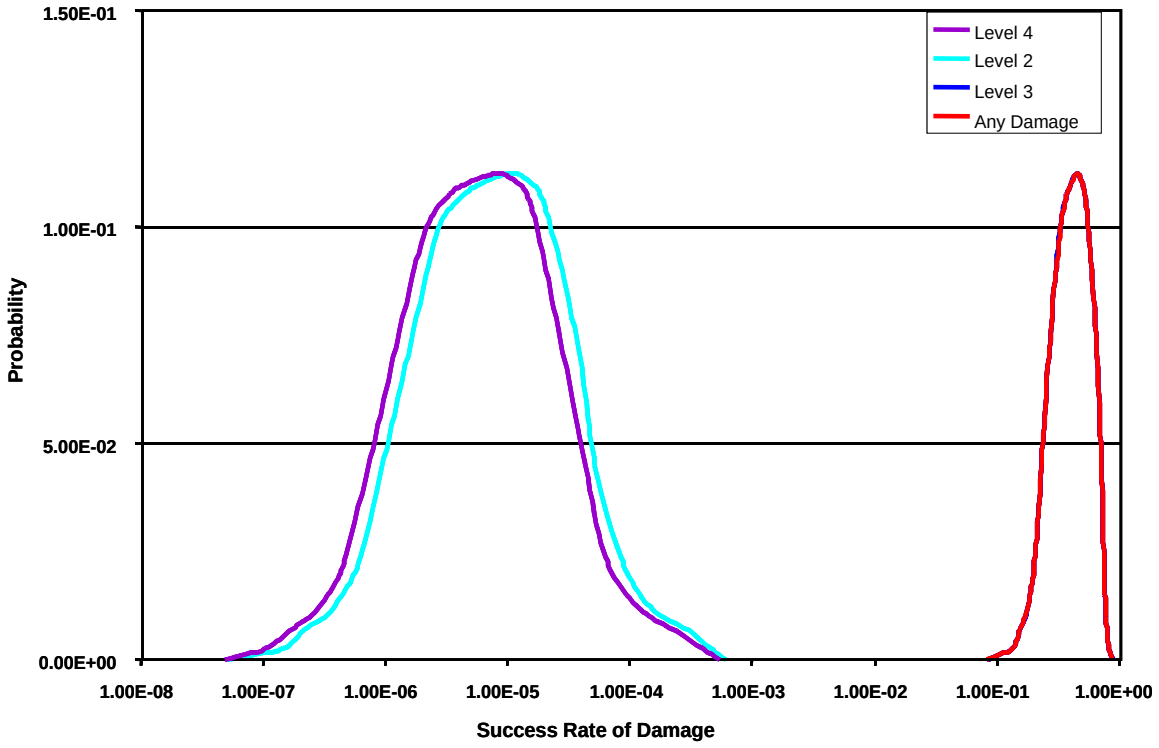

Figure 4-7. Combined Results for All Damage Levels

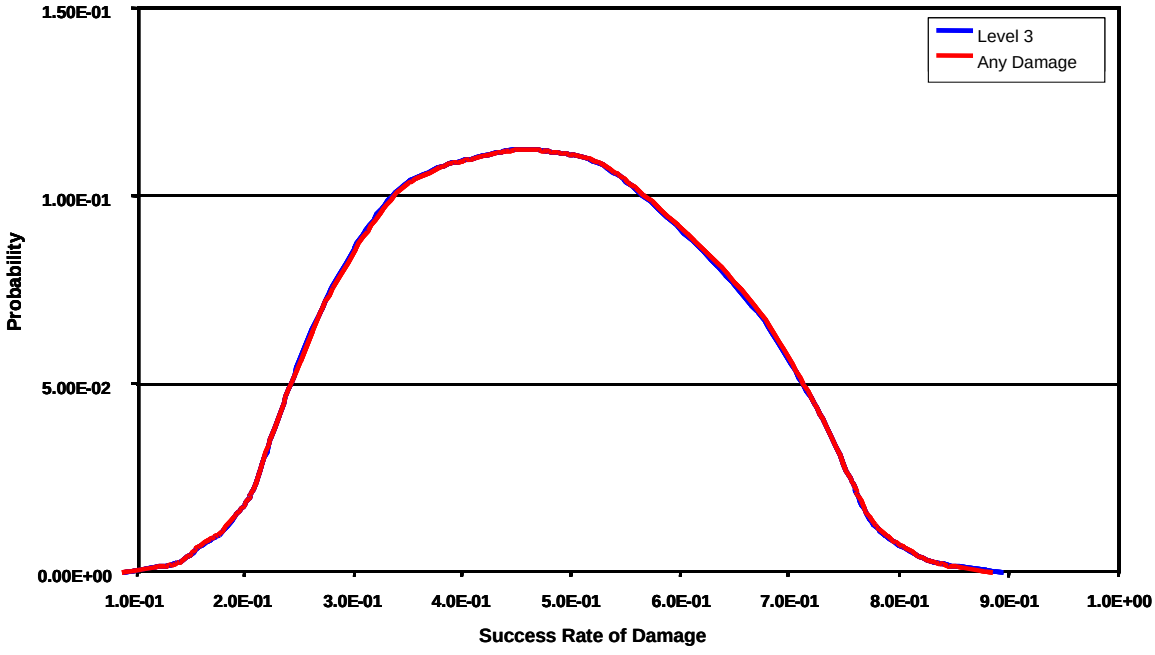

Figure 4-8. Results for Damage Level 3 and Any Damage

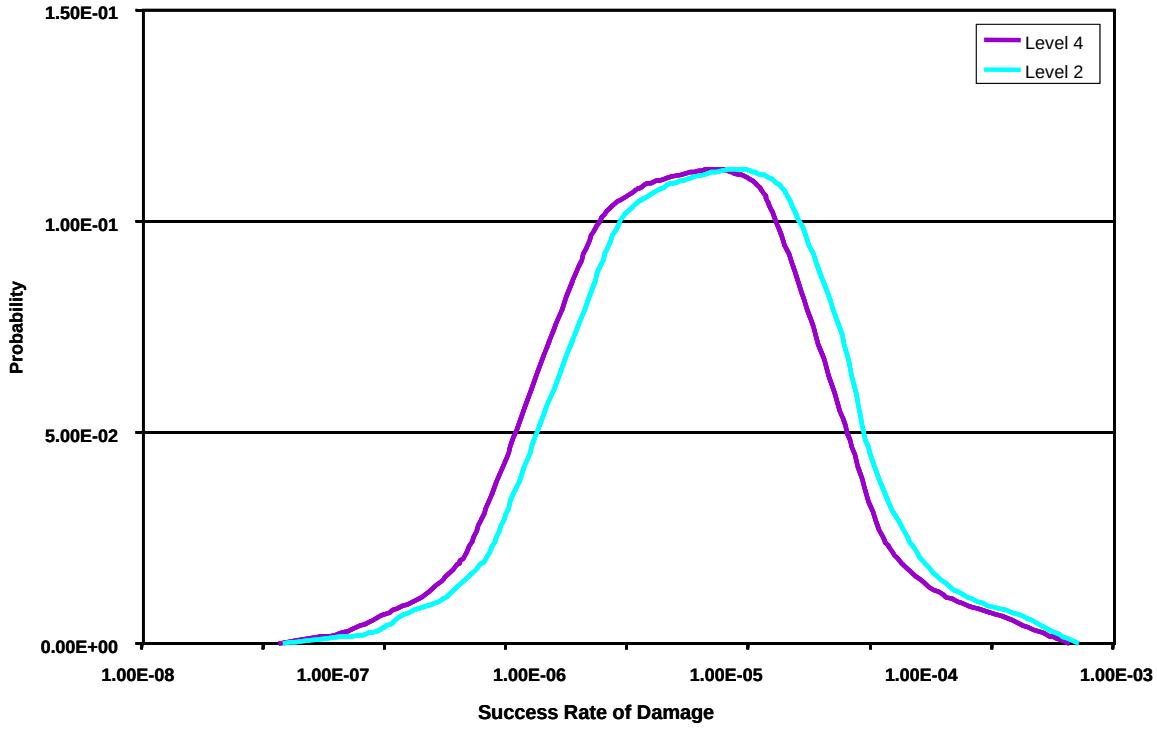

Figure 4-9. Results for Damage Level 2 and Damage Level 4

Table 4-5 summarizes selected parameters of the uncertainty distribution for each level of damage. It is evident from Figures 4-7 through 4-9 and from Table 4-5 that, for the model discussed, the likelihood of a successful attack is much greater for a physical attack (any damage; Damage Level 3) than for a cyberattack (Damage Levels 2 and 4).

Table 4-5. Selected Parameters of Uncertainty Distribution for Each Level of Damage

Success Rate of Attack				
Damage Level	**5th Percentile**	**Median**	**95th Percentile**	**Mean**
Any Damage	2.1×10^{-1}	4.2×10^{-1}	7.4×10^{-1}	4.4×10^{-1}
Level 2	6.8×10^{-7}	6.9×10^{-6}	5.9×10^{-5}	2.0×10^{-5}
Level 3	2.1×10^{-1}	4.2×10^{-1}	7.4×10^{-1}	4.4×10^{-1}
Level 4	5.2×10^{-7}	5.7×10^{-6}	5.3×10^{-5}	1.6×10^{-5}

Step 6. Interpret the results to guide the risk-management process.

INTERPRET THE RESULTS

We must now ask if the methodology described in Chapters 2 and 3 meet our expectations. To focus our answer, we must first revisit the questions that were the basis for proposing a risk-based methodology. The methodology is intended to answer such questions as what are the threats and vulnerabilities; what are the contributing factors, and how do they rank in importance; what actions will have the biggest payoff in terms of risk reduction for the amount of resources invested.

Threats and Vulnerabilities

The two modes of terrorist attacks in the example were a physical damage and a cyberattack. The elements of the electrical grid considered in the vulnerability assessment were substations, transmission lines, and the SCADA systems. The interface between the threat assessment and the vulnerability assessment is the actual attack on the grid itself. This initiating event (i.e., the nature of the attack) is the output of the threat assessment and the input for the vulnerability assessment.

Events and activities leading up to the attack are part of the threat assessment. For the cyberattack, the assessment included many of the events involved in gaining access to SCADA systems for the attackers to be in a position to initiate commands that actually result in grid damage. To that extent, the sample application involved some aspects of a threat assessment, even though it did not go all the way back to the point of the terrorists' decision to launch an attack, which would have required intelligence information and an extensive information search. For the physical attack, only the vulnerability of the grid was considered. The attack modes were hypothesized, but the scenarios for accessing the SCADA system demonstrate the type of scenario structuring required for a comprehensive threat assessment.

The methodology revealed that the grid would be much more vulnerable to physical attacks than to cyberattacks, but the uncertainties associated with cyberattacks would be greater. The attackers would successfully cause a long-term outage in Network 1, short-term outages throughout the region, or both approximately 10 times in every 23 attempts. The analysis clearly shows that the attackers would have the highest likelihood of causing long-term power outages in Network 1. Thus, the sample application demonstrates that even an abbreviated risk assessment can yield meaningful results on the vulnerability of the grid and the threat of a cyberattack.

Contributing Factors

According to the analysis of a cyberattack in the example, there would be a very low likelihood of successful intrusion into the regional SCADA control systems, although there was a great deal of uncertainty in the estimates. For example, there was 90 percent confidence that the success rate for a transient disruption of regional power would be between 6.8×10^{-7} and 5.9×10^{-5}, or approximately 1 success in every 1.5 million to 17,000 attempts. However, if the attackers were successful in disrupting regional power, it was quite likely that they would also cause long-term damage to Network 1 (i.e., Damage Level 4 was only slightly less likely than Damage Level 2).

In this example, even though cyber-initiated events did not constitute a major threat, they could not be ignored. First, there was a wide range of uncertainty in the assessment. Second, unlike a physical attack in which the risks of repeated attempts to the terrorists would be high, a cyber-initiated interruption of power could be attempted many times with very little investment and very little risk to the terrorists. The analysis revealed a need to develop more information to reduce the uncertainty and to explore ways of discouraging repeated attempts.

Coordinated Physical Attacks on System Hardware

In this example the focus was on Network 1, which was determined to be the most vulnerable to long-term outages. The overall vulnerability of Network 1 was most strongly determined by the relatively high vulnerability of substation S2, in spite of the fact that this substation was not as important to the network's electrical stability as the other more secure substations.

Table 4-6 shows that successful attacks on substation S1 would contribute to approximately 69 percent of Damage Level 3; successful attacks on substation S2 would contribute to approximately 96 percent of Damage Level 3; successful attacks on substation S3 would contribute to approximately 62 percent of Damage Level 3. (These so-called fractional-importance measures are simply the sum of scenarios that include damage to each substation, divided by the total number of scenarios.) Therefore, the overall vulnerability of Network 1 is most strongly determined by the relatively high vulnerability of substation S2, even though this substation was not individually as important to the network power generation and transmission interties as the more secure substation S1.

Table 4-6 summarizes the results of an assessment of coordinated physical attacks on the system hardware.

A full-scope risk assessment of a real electrical grid would take the contributing-factor question to a much more detailed level than was possible in this example. A full-scope risk assessment would also consider the next level of consequences (i.e., injuries and fatalities to workers and the public). The principles of the analysis, however, would be the same.

Actions with the Greatest Payoff

To avert cyber-initiated attacks, steps could be taken to reduce the uncertainties in the analysis and to find ways to discourage repeated attempts. For coordinated physical attacks, one very clear action to consider would be to improve the security of substation S2, which was identified as the principal contributor to long-term outages for Network 1. This priority might not have been evident without an integrated assessment of the vulnerabilities and the potential consequences of the failure of each substation. It is also very clear from Table 4-6 that attacks on multiple substations would greatly increase the likelihood of Network 1 failure. Thus, substation security in general would be an important consideration in improving the security of the regional grid. The relative importance of the subsystems to overall vulnerability would not be readily apparent without an integrated model that systematically evaluated each contribution to damage.

Table 4-6. Risks of Coordinated Physical Attack

Damage to Substations	Likelihood of Success*	Fraction of Total Damage Level 3
S1 and S2 and S3	1.74×10^{-1}	39.4%
S1 and S2	1.17×10^{-1}	26.5%
S2 and S3	8.72×10^{-2}	19.8%
S2	4.42×10^{-2}	10.0%
S1 and S3	1.02×10^{-2}	2.3%
S1	5.20×10^{-3}	1.2%
S3	3.61×10^{-3}	0.8%

* Combined likelihood of successful substation attack and failure of Network 1 as a consequence of the substation damage.

Once developed, models become key elements in a systematic risk management process to evaluate the effectiveness of proposed improvements. The updated analysis results display the corresponding changes to the overall grid risk profile and reorder the contributors to each damage level. Based on these risk-based insights, systematic examination of successive improvements would continue until an acceptable level of overall risk was achieved.

CONCLUDING REMARKS

This example is intended to illustrate how QRA can be used to "turn up the microscope" to expose the risk of an event that is either catastrophic or could become catastrophic. No extensive analysis is necessary in situations where the risks are apparent (i.e., when the threats and vulnerabilities can be easily identified). Obvious steps can be taken to reduce the vulnerability to a terrorist attack of many important assets in conventional facilities and buildings. Risk reduction in those situations may include improving ventilation systems, emergency action training, improving security, providing rapid escape systems, identifying protective staging locations, and upgrading emergency response capabilities.

CONCLUSION AND RECOMMENDATION

Conclusion. Quantitative risk assessment is an effective method of exposing the risks of complex systems to events that could lead to catastrophic consequences. The hallmark of a

quantitative risk assessment is the quantification of uncertainty—uncertainty is the risk of greatest concern.

Recommendation. Quantitative risk assessment should be applied in cases where the consequences can be catastrophic and where there is great uncertainty about the risk scenarios and contributing factors. Meanwhile, the government and private sector should act quickly to reduce the risk to those assets where the payoff can be readily determined.

REFERENCES

Alvey, J. 2002. Digital terrorism: holes in the firewall? Public Utilities Fortnightly 140(6).

Amin, M. 2002. Security challenges for the electricity infrastructure. Security and Privacy35(4): 8-10.

DOE (Department of Energy). 2002. National Transmission Grid Study. Washington, D.C.: U.S. Department of Energy.

EPRI (Electric Power Research Institute). 2002. Electricity Infrastructure Security Assessment: EPRI Security Overview. Palo Alto, Calif.: Electronic Power Research Institute (distribution limited).

NAERC (North American Electric Reliability Council). 2002. The Electricity Sector Response to the Critical Infrastructure Protection Challenge. Princeton, N.J.: North American Electronic Reliability Council.

CHAPTER 5

THE INFORMATION FOUNDATION FOR QRA

In this chapter, several issues are addressed associated with information pertinent to risk assessment and the institutional environment from which the necessary information originates. First, the new U.S. Department of Homeland Security is responsible for establishing organizational arrangements, roles, and processes to protect Americans and their infrastructure. The department will serve as matrix manager of the activities of its agencies and as coordinator of interagency activities. The key to good decisions to counter terrorism will be organizational structures that support the management of reliable, timely information. Given the complex issues involved, the Department of Homeland Security should place the highest priority on the effective collection, fusion, analysis, and sharing of relevant data. Dozens of federal departments and agencies and hundreds of local agencies have extensive databases that could be mined. The involvement of private-sector organizations, which will be essential, can be facilitated through consortia and other public/private collaborations, such as information sharing and analysis centers (ISACs).

The Department of Homeland Security should also focus on two areas: (1) integrating and analyzing information from all sources; and (2) ensuring that mechanisms for assessing and mitigating threats are in place. A DHS priority should be stimulating interest in sector-by-sector, structured risk assessments and related risk-reduction activities by federal and local governments, private-sector owners of targets, law enforcement organizations, and first responders. The objective must be to ensure that mechanisms to address threats and vulnerabilities to terrorism are in place throughout the country. Regulations, government guidelines, voluntary worldwide manufacturing standards by the International Standards Organization (e.g., ISO 9000), and other mechanisms could be used to help achieve this goal.

THE INFORMATION DIMENSION

Risk mitigation achieved through threat and vulnerability assessments is built on the information and evidence base that forms the foundation for all risk assessments. The better the data, the better the assessment and the better the decisions. There are obviously specific challenges: enabling connectivity among government organizations; balancing civil liberties; the identification and linking of existing databases; the involvement of risk experts in assessing data quality; and greater use of existing technologies in information management. In the risk assessment process, the gatherers and end-users of information will most likely be very different bodies, crossing geographical, government, nongovernment, intelligence classification, and even political boundaries. The infrastructure for coordinating, sharing, accessing, disseminating, desensitizing, declassifying, and ultimately trafficking information is essential to risk assessment.

One of the themes of this report is that risk assessments of terrorist attacks should be based on supporting evidence for that attack. The term "evidence" is purposely chosen over the often-used term "data" because the input to a risk assessment is not simply data. All three terms (evidence, data, and information) are often used interchangeably, they tend to have different meanings based on context. "Data" often carries with it the connotation of statistics, measurements, failure rates, etc., for the purpose of making calculations. "Information" is usually interpreted more generally as having to do with communicating or receiving knowledge or intelligence; information is the term primarily used in this chapter. Input to threat assessment, for example, is much easier to comprehend in terms of information and intelligence, than in terms of occurrence rates. "Evidence" is used primarily in the legal field and has the connotation of something that furnishes proof.

The Department of Homeland Security and the intelligence agencies face an enormous problem—information fusion (see Box 5-1). Tens of millions of terrorism-related information bits and pieces with different characteristics are housed in thousands of databanks both inside and outside the government that serve a variety of end-users. Some of the information is dated; some is of questionable reliability; and some is just plain wrong. Some of the databanks are disorganized; some are not even recognizable as relevant to terrorism; and some are accessible only to authorized users with unique computer skills and special clearances. Internationally, the situation is even more complicated. Not only are data processed and stored using a variety of database schema, but information itself is often described in different written languages and has meanings based on cultural interpretations. Blending foreign and domestic intelligence is a daunting task, especially when raw intelligence or details of how or where intelligence was gathered are involved. At least initially, much of DHS's intelligence fusion and coordination activity will depend on processed intelligence, rather than raw intelligence, which could affect the objectivity of the results (Senate Report 107-63).

Box 5-1. Information Fusion

Risk assessment greatly benefits from high-quality information that is likely to come from a variety of sources, including local and state agencies, satellite images, embedded sensors, video systems, and human intelligence. Information fusion then comes into play for processing this qualitative and quantitative data in such a way as to provide an objective picture of current reality, generally on a visual basis.

Research in this area seeks to provide new techniques for multi-attribute decision making under conditions of uncertainty through the use of models involving a choice of attributes, prediction of expected values for those attributes, and algorithms to provide decision alternatives. Much of this research also involves the determination of preferences based on utility functions so that these methods can be applied, hopefully leading to an optimal choice.

On a practical level, much of this work is done by computer using mathematical techniques developed for topology, optimization, linear and dynamic programming, fuzzy logic, neural networks, and Bayesian analysis. Classic examples of applications are managing post-event activities, such as evacuations of large numbers of people, routing of emergency services, the deployment of defense forces, and the assessments of actual physical damage. Applied to counterterrorism, information fusion can be used to synthesize quantitative and qualitative data so that a picture may emerge of likely threat scenarios. However, this work is still somewhat in its infancy and tends to be highly interdisciplinary but promises to provide enhanced capabilities for the risk assessment community.

Identifying and Linking Existing Databases

One approach to bringing all data sources together is through the Terrorist Threat Integration Center (TTIC) (Box 5-2). The TTIC will need to compile an up-to-date catalogue of existing databases with descriptions of their contents and organize the catalogue in a way that will render the information useful to all users. One organizing principle of the catalogues should

be the relevance of databases to analyzing terrorism scenarios and conducting threat and vulnerability assessments. Separate catalogues may be necessary for identifying classified and unclassified databases, publicly available and commercial databases, and on-line and limited-access databases.

Box 5-2. The Terrorist Threat Integration Center

In his State of the Union Address on January 28, 2003, President Bush announced that he had instructed the director of the CIA and the director of the FBI, working with the attorney general and the secretaries of homeland security and defense, to develop the first unified Terrorist Threat Integration Center. This new center will merge and analyze terrorist-related information collected domestically and abroad. The center, as envisioned, will oversee national counterterrorism activities and maintain shared databases as well as an up-to-date database of known and suspected terrorists that will be accessible to federal and nonfederal officials and entities. Ideally, the center will have access to all intelligence information—from raw reports to finished analytic assessments—available to the U.S. government. The new initiative requires the U.S. Department of Homeland Security's Information Analysis and Infrastructure Protection Directorate to add new capabilities in the area of information analysis and infrastructure protection.

The directorate will:

- Perform comprehensive vulnerability assessments of the nation's critical infrastructure and key assets.
- Receive and analyze terrorism-related information from the Terrorism Threat Integration Center, as well as open sources, the public, private industry, state and local law enforcement, and the entire federal family.
- Map the threats against vulnerabilities, to develop a comprehensive picture of the terrorist threat and our ability to withstand it.
- Take and facilitate action to protect against identified threats, remedy vulnerabilities, and preempt and disrupt terrorist threats, as consistent with the operational authorities of the department's constituent agencies.
- Take a lead role in issuing warnings, threat advisories, and recommended response measures to public safety agencies, elected officials, industry, and the public.

Source: White House Fact Sheet, 2003.

Linking databases that were not designed to be linked is difficult and often requires the development of middleware to enable data exchange. In the design of new databases, connectivity must be a design consideration, and the DHS will need a comprehensive understanding of new and old databases as well as expertise in the theory and logistics of deploying diverse, but interconnected database and knowledge systems to assist agencies in data sharing. DHS will also need expertise in the creation of physical communication, authorization, and control channels between interagency and intra-agency databases, as well as theoretical and practical experience in knowledge engineering. Database linkage should facilitate the identification of duplicate entries, that is recurring places, individuals, and objects that appear in multiple databases, as well as elicit the presence of more complicated relationships, such as patterns in one database that predict a pattern in another database.

Enhancing the Quality of Data

The successful integration of disparate pieces of information into a coherent risk assessment requires being able to assess the validity or truth of the data elements used in the assessment. Effective linking of information systems will require a data-validity standard that can be applied across numerous databases. If a validity scoring system is already in place, it can either be adopted or adapted. Connecting information to form a picture with limited uncertainty requires that data elements be accurate, reliable, and timely. However, because data sets are often incomplete and many key data fields may be missing, estimates to fill a "data hole" will be necessary. Sometimes, the missing field may be available within an agency or in another agency's database that is available for sharing. In these cases, the data hole can be plugged either by direct substitution of the data value from the other database or by inserting a "pointer" to the correct value in a shared database.

In some cases, no value is available for a key field, or even for an entire record. In these cases, the data values are often imputed either by a mathematical or statistical algorithm or by probabilistic methods. To appreciate the level of uncertainty, it is crucial that the user of these data be informed which values are imputed and which data points are not imputed. It is often much easier to impute estimates of the data holes at the time the data is first imputed when the characteristics of the data and their accuracy are known. Thus, estimates of missing information should be included but should be labeled as imputed and should include references to the method of imputation.

A number of technologies and statistical tools are relevant to address these information challenges. These should be incorporated, where possible, when connecting databases. These include data mining applications (algorithmic processing to discern general patterns in a large volume of specific cases); data integration methods (combining data collected from multiple sources with different sampling rates or data schemas; see the discussion on Bayesian inference in Chapter 3); language processing (managing large volumes of text and speech, often recorded in different languages); image and video processing (face recognition, automatic recognition, and classification of biological spore images); and evidence combination techniques (combining information from multiple sources to reach a conclusion) (NRC, 2002).

INFORMATION CHALLENGES SPECIFIC TO THREAT ASSESSMENTS

Assessing threats, that is, detecting the prospects of a specific attack during the planning or early execution phase, is analogous to finding a needle in a haystack. But a catastrophic attack requires extensive preparations, which may generate small bits of evidence from disparate sources that could be pieced together by analysts. The sources of evidence may not be only the usual information and intelligence agencies, but also security organizations in the business of searching for terrorism information from a wide variety of sources, such as "search" perturbations on the Internet. It is with respect to catastrophic attacks that the principles and practices of quantitative risk assessment have their greatest value. The structuring of catastrophic attack scenarios could be one of the most important short-term benefits of quantitative risk assessment. Mining current information sources for the possibility of catastrophic attacks and having access to intelligence experts as well as knowledgeable scientists and engineers could have important near-term benefits in combating terrorism.

Disinformation, Misinformation, and the Timeliness of Data

For threat assessments, the quality of data can be compromised by the terrorist's desire to obscure "signals" that might be monitored by the intelligence community. Unlike vulnerability assessments, where risk and failure characteristics can be defined with a greater degree of certainty, malevolent adversaries have clear incentives to impede the monitoring of their activities. Thus, terrorists may cover up indicators of their plans, which could lead to underestimates of risk. The terrorist may spread disinformation to essentially create a "false positive" leading to the overestimation of alternative scenarios. Finally, the terrorist might spread misinformation, leaking intentionally false risk-sensitive information that could change the dynamics on how monitoring is conducted. The timeliness of data is also crucial. External events or changes in terrorist resources can easily change the likelihood of a terrorist favoring

one scenario over another. External events could cause large swings in the corresponding threat assessments.

INFORMATION CHALLENGES SPECIFIC TO VULNERABILITY ASSESSMENTS

An accurate vulnerability assessment is influenced by (1) who does the assessment and (2) the quality of the available information. In many cases, private sector organizations may be best suited to perform a vulnerability assessment, which creates unique challenges on whether and how to provide potentially sensitive information on a variety of terrorist motives and capabilities. This problem raises a host of issues associated with security clearances, document classification, and private-sector proprietary information.

Barriers to Information Sharing

Factors that inhibit information sharing between government and the private sector and within the private sector include privacy concerns, competitive pressures, a lack of perceived need, and concerns about revealing legally actionable flaws, shortcomings, or vulnerabilities. Between companies, antitrust concerns include price fixing, restraint of trade, systematic exclusion of competitors, and discrimination against certain customers. For example, to reduce vulnerability to attack, companies could agree to improve their capabilities through the deployment of expensive technology and pass costs on to customers. Such a deployment and charging mechanism might have to overcome antitrust hurdles.

The Freedom of Information Act (FOIA) might also limit information sharing. FOIA gives the public broad access to government records and databases, except for information on national security, trade secrets, and confidential business matters. An unfortunate result of FOIA has been that the private sector has resisted sharing information with government because of potential public disclosure. Even specific, agreed-upon restrictions to disclosure are not always passed from one government department to another; an uninformed department may release information, or a careless government employee may release sensitive information. Under FOIA, government entities and the private sector can negotiate a Memorandum of Understanding to restrict public disclosure, but the effectiveness of these agreements has not been thoroughly tested. Many companies worry that future court rulings and interpretations could result in the release of previously restricted information. The President's Commission on Critical Infrastructure Protection called for a new FOIA exemption to "encourage and protect" information on critical infrastructures. Opponents to the changes argue that existing exemptions are adequate.

Concerns over release by the federal government of information relating to critical infrastructure vulnerabilities have been addressed, to some degree, by the Critical Infrastructure Information Act of 2002 (P.L.107-296). Section 214 provides that critical infrastructure information voluntarily submitted to the Department of Homeland Security shall be exempt from distribution under FOIA as long as an express statement on dissemination is included with such information to the effect that such information is voluntarily submitted and protection of the U.S.A. Patriot Act is sought. However, one potential loophole in the law is that it exempts information submitted to the Department of Homeland Security but not to other federal agencies. Moreover, the scope of the information protected is unclear and arguably extremely broad. For

instance, the term "critical infrastructure" as defined by the USA Patriot Act [Section 106(e) of P.L. 107-56] includes both physical and virtual systems and assets. Finally, independently obtained information held by the federal government or the Department of Homeland Security is not exempt from FOIA provisions under the Patriot Act.

Sharing Information with Local Governments and the Private Sector

In the event of an attack, local governments will be the first responders to contain the damage and provide emergency services. Local governments must be prepared to determine the nature of the attack and respond with emergency resources, including firefighters, police, ambulances, and hospital care. In addition, they must be ready to contain the attack by isolating the area and protecting other vulnerable targets. The private sector will be called upon to isolate and protect private targets and assist in protecting public targets. Currently, neither first responders nor the private sector nor the public has the ability to access and interpret intelligence information and related threat assessments so they can assess their own vulnerabilities.

Many individuals and institutions at different levels of government and society are now involved in making vulnerability assessments and consequence-management decisions. But much of the intelligence data is classified and therefore not available or has been processed and watered down to the point of being almost useless. Because of concerns about leaks, intelligence organizations are likely to involve others only when they have reasonably reliable information that specific areas or facilities may be immediate targets. Sharing classified information with state and local governments raises many concerns about clearances and decisions about who has a need-to-know.

The private sector will require incentives, prodding, and support to recognize and respond to the risks to private sector targets. The private sector is complex, and heterogeneous, ranging from individual enterprises to giant mega-corporations having enormous assets, and how one characterizes vulnerabilities is largely dependent on the type of organization characterizing it. Each company knows itself and its vulnerabilities, and such knowledge is central to the security risks that need to be considered. Although improved security can be expensive, the federal government can take steps to keep costs down by encouraging and rewarding the introduction of industry standards and good practices. Government incentives might include: tax incentives for measures that minimize damage from terrorism; limitations on liability; tax deductions for selected private-sector security investments; adjustments of antitrust laws that inhibit intra-industry and inter-industry cooperation; and government awards for industry initiatives. Counterterrorism programs sponsored by professional societies and trade associations and initiatives at the local level to strengthen linkages between private-sector facility managers and law enforcement organizations can be effective. Public-private partnerships such as Operation Safe Commerce, a Transportation Security Administration initiative designed to improve the security of international maritime container shipments, can also engage the private sector in improving national security.

Overcoming Barriers

Industries and companies that already share information related to terrorism have established some good role models. A good example is the telecommunications industry, which began sharing information on the vulnerabilities of public switched networks to Russian

terrorists during the Cold War (NRC, 1989). In 1963, President Kennedy established the National Communications System (NCS), an interagency organization to ensure the reliability and availability of national security and emergency preparedness communications. In 1982, President Reagan created a government/industry committee, the National Security Telecommunications Advisory Committee (NSTAC), to share information on the vulnerabilities of the public switched network infrastructure.

The primary "customer" for work performed by NSTAC is the NCS, which manages the National Coordinating Center, a government-industry partnership created in 1984 to coordinate the response to disruptions in the NCS. An ISAC for the telecommunications industry was established through the National Coordinating Center in January 2000.

Another means of information sharing in the telecommunications industry is through the Federal Communication Commission (FCC), which requires that failures in company networks that are part of the overall public switched network be reported. The FCC compiles a summary of all reported failures, which are available to all companies in the network.

The National Research Council of the National Academies provides another means of information sharing and vulnerability assessment. NCS has sponsored studies by the NRC on vulnerability assessments, such as *Growing Vulnerability of the Public Switched Networks*, conducted by a committee of industry experts (NRC, 1989).

BUILDING THE FOUNDATION

Two critically important issues will have to be overcome for successful quantitative risk assessments to combat terrorism. They are: (1) gaining access to information that already exists, and (2) recognizing that limited data and the attendant uncertainties of catastrophic consequences were the background against which quantitative risk assessment was developed. Limited information is not a legitimate reason for not moving forward with analyses of the risk of terrorism. Experience has indicated that the best course for the development of effective databases to support risk assessment is through the application of such analyses and the explicit exposure of information weaknesses that can then be corrected.

Enough sources of information are available to support meaningful analyses when and if the institutional problems of sharing, accessing, and quality checking can be resolved. Many precedents have been established for linking disparate sources of information through computer conduits for purposes of fact finding and analysis.

The information base for vulnerability assessment is much more developed and accessible than it is for threat assessment. This is because: (1) formal vulnerability assessment is reasonably well developed and has been practiced in many industries for several decades, and (2) threat assessment information is dispersed, more guarded, and more limited.

Linking threat scenarios with vulnerability scenarios is the key to the development of integrated terrorism risk models. Even though the information on threats will most likely be the

greatest source of uncertainty in a fully integrated (threats and vulnerabilities) risk assessment of any terrorist attack, it is important that the two be linked.

The effective use of current information sources requires specialization, not only to support risk assessments, but also for end-users (including first responders), facility owners and operators, private citizens, institutions, and risk managers. Processors and interfaces should be implemented to sort and customize information for many different users.

Classification of information should not be a deterrent to having the very best information available to perform quantitative risk assessments of potential terrorist attacks. It is likely that most, if not all, of the risk assessments will have to be classified to avoid blueprinting candidate terrorist attack scenarios.

CONCLUSIONS AND RECOMMENDATIONS

Conclusion 5-1. The problem of information fusion is enormous. Data are housed in thousands of databanks inside and outside the government to serve a variety of end-users. There are serious issues of accessibility, relevancy, organization, reliability, and processing of data for interpretation.

Recommendation 5-1. The U.S. Department of Homeland Security should place the highest priority on the effective collection, fusion, and sharing of relevant data. The involvement of private-sector organizations will be essential, and consortia and other collaborative mechanisms, such as information sharing and analysis centers (ISACs), should be used whenever possible.

Conclusion 5-2. Coordinating government databases will require improved interfaces and standardization of future data. Interfacing of data collected by different organizations for different purposes may not be possible, however, which means some existing data probably will be of little use.

Recommendation 5-2. The U.S. Department of Homeland Security should continue to pursue aggressively the task of identifying and prioritizing possible attack scenarios for developing and implementing defensive actions. The department must have strong in-house capabilities to ensure the compatibility, quality, and accessibility of information. Working with many federal agencies and end-users throughout the country, the department should continue its efforts to find disparate databases with data relevant to countering terrorism and to integrate them into the larger collection of data.

Conclusion 5-3. Many of the assets targeted by terrorists are owned by the private sector, which must take steps to reduce their vulnerabilities.

Recommendation 5-3. Industry associations and companies should share information on vulnerabilities and take steps to reduce them. Where possible, industry sectors should establish information sharing and analysis centers.

Conclusion 5-4. The defense against terrorism must be addressed not only at the national level, but also at the state, local, and private-sector levels. Nevertheless, much of the critical infrastructure is in the private sector, thus requiring close coordination with actions by the local and regional governments that provide police, emergency first responders, and other support services. This coordination must be made more seamless.

Recommendation 5-4. The U.S. Department of Homeland Security should be stimulating interest in sector-by-sector, structured risk assessments and related risk-reduction activities by federal and local governments, private-sector owners of targets, law enforcement organizations, and first responders.

REFERENCES

NRC (National Research Council). 1989. Growing Vulnerability of the Public Switched Networks. Washington, DC. National Academy Press.

NRC. 2002. Making the Nation Safer: The Role of Science and Technology in Countering Terrorism. Washington, DC. National Academy Press.

White House. 2003. Strengthening Intelligence to Better Protect America. Fact Sheet, January 28, 2003. Washington, D.C.: White House.

APPENDIX A

HISTORICAL PERSPECTIVE OF QUANTITATIVE RISK ASSESSMENT

Risk assessment, the critical building block of risk management, has been part of decision analysis since man was able to reason. A caveman hunting for food had to decide whether to get close to a beast before throwing his spear, thus accepting the risk of being attacked himself, or to trust his throwing skills from a greater but less risky distance. But the formalized process of making decisions about risks began much later. Probability theory, the foundation of contemporary risk analysis, was based on discoveries in the sixteenth and seventeenth centuries by notable scholars, such as Girolamo Cardano, Galileo Galilei, Blaise Pascal, Pierre de Fermat, and Chevalier de Méré.

Cardano and Galileo made important contributions in the 1500s on how to express probabilities and frequencies of past events, Pascal contributed to concepts of decision theory and statistical inference in the mid-1600s, and Fermet and de Méré made major contributions to the theory of numbers about the same time. Other major contributors during the seventeenth century were Christen Huygens, who published a popular textbook on probability theory, Gottfried Wilhelm von Leibniz, who suggested applying probability methods to legal problems, and members of a Paris monastery named Port Royal. The Port Royal group produced a pioneering work of philosophy and, probably, the first definition of risk: "Fear of harm ought to be proportional not merely to the gravity of the harm, but also to the probability of the event." Two excellent sources on the history of the risk sciences are by Bernstein (1996) and Rechard (1999).

Many consider Thomas Bayes, an English minister, the real father of contemporary risk assessment. In the mid-1700s, he developed a theorem rooted in fundamental logic for combining old information with new information for the assignment of probabilities. Bayes Theorem, followed by the publication in 1812 of *Théorie analytique des probabilities* by the French mathematician Marquis Pierre Simon de Laplace, provided the primary basis of contemporary probability theory. Diverse problems, such as gambling strategies, military strategies, determining mortality rates, and debating the existence of God, were the subjects of early analytical explorations and precursors to the new science of risk assessment. Among the scholars who contributed to risk assessment during the twentieth century are Harold Jeffreys (1957), Howard Raiffa (1996), and E.T. Jaynes (2003).

The widespread, formal application of risk assessment to critical infrastructure began in earnest in the late 1900s. Applications in the insurance and financial fields were more statistical (actuarial) than probabilistic, more experience-based than analytical, more qualitative than quantitative. Only when societies began depending more on technological systems involving large inventories of hazardous materials did investigators begin to look for more scientifically based ways to assess risks. The particular need was for a method of assessing the likelihood of catastrophic events that could do great harm to public health and the environment.

A NEW WAY TO THINK ABOUT SAFETY ASSESSMENT

In the past four decades, a great many analysis methods have been developed and put into practice to support scientifically based risk assessments and decision analyses. These methods have had a major impact on public policy in diverse areas, including public health, environmental regulation, safety regulation, and the performance of technological systems, especially systems involving hazardous materials. Pioneering studies by the nuclear industry in the 1960s led to significant improvements in the effectiveness and sophistication of safety analyses. F.R. Farmer of the United Kingdom proposed a new approach to nuclear power plant safety based on the reliability of consequence-limiting equipment (Farmer, 1964). A series of studies was performed by Holmes and Narver, Inc., a U.S. engineering firm under contract to the then U.S. Atomic Energy Commission. The final report in the series advocated, with examples, the need for much greater use of advanced systems-engineering methods of modeling the reliability of safety systems. The authors made explicit reference to the use of logic tools, such as fault-tree methodology, which has its roots in "switching theory" developed by the telecommunications field (Holmes and Narver, Inc., 1967). At about the same time, a Ph.D. thesis was published that proposed a methodology for probabilistic, integrated systems analysis for analyzing the safety of nuclear power plants (Garrick, 1968). All of these works focused on the engineering side of the risk issue. Others were thinking much more globally. For example, Chauncey Starr published a seminal paper on links between societal benefits and technological risk that provided a bridge between technological risk and critical social issues. Starr addressed the difficult issue of voluntary versus involuntary risk (Starr, 1969).

THE BREAKTHROUGH

The breakthrough in probabilistic risk assessment of technological systems came in 1975 with the publication of the *Reactor Safety Study* by the U.S. Atomic Energy Commission under the direction of Professor N.C. Rasmussen of the Massachusetts Institute of Technology (USNRC, 1975). This project, which took three years to complete, marked a turning point in the way people think about the safety of complex facilities and systems. The *Reactor Safety Study* proposed a wide range of applications for risk assessment, not only for nuclear power plants and other technological systems (e.g., chemical and petroleum facilities, transportation systems, and defense systems), but also for environmental protection, health care, and food safety.

The *Reactor Safety Study* inspired many first-of-a-kind risk assessments in industry that led to major advancements in the application of quantitative risk assessment. One important example was the probabilistic risk assessments of the Zion and Indian Point nuclear power plants sponsored by the owners and operators of the plants. New methods were introduced in those assessments that have become standards of many quantitative risk assessment applications (PLG et al., 1981, 1982). The methods included the treatment of uncertainty, a framework of risk assessment embedded in the triplet definition of risk (Kaplan and Garrick, 1981), common-cause failure analysis, importance ranking of risk contributors, models for calculating source terms, and improved models for calculating off-site health effects.

The U.S. Nuclear Regulatory Commission issued an update of the original *Reactor Safety Study* (USNRC, 1990). This report was based on a review of five nuclear power plants and emphasized severe accidents and containment performance. The safety risks of nuclear power plants were calculated to be less than in the original study, primarily because of improvements in

computational methods and a better understanding of accident phenomena. Finally, a number of high-level, comprehensive reviews of major risk assessment programs have made important contributions to quantitative risk assessment. Three have to do with nuclear, space, and chemical weapons disposal (Apostolakis, et al., 2002, Apostolakis, et al., 1996, and Kastenberg et al., 1988).

THE STEPS OF A QUANTITATIVE RISK ASSESSMENT

Although the scope, depth, and applications of quantitative risk assessments vary widely, they all follow the same steps:

1. Define the system being analyzed in terms of what constitutes normal operation, and points of vulnerability to serve as a baseline reference point.
2. Identify and characterize the sources of danger, that is, the hazards (e.g., stored energy, toxic substances, hazardous materials, acts of nature, sabotage, terrorism, equipment failure, combination of each, etc.).
3. Develop terrorist attack scenarios to establish levels of damage and consequences.
4. Adopt risk metrics that reflect the likelihoods of different attack scenarios in terms of target and collateral damage, and quantify the scenarios based on the totality of relevant evidence.
5. Assemble the scenarios according to damage levels, and cast the results into the appropriate risk curves and risk priorities.
6. Interpret the results to guide the risk-management process.

These steps provide the answers to the three fundamental questions of risk (the "triplet definition"): what can go wrong, how likely it is to go wrong, and what the consequences will be.

APPLICATIONS

Risk assessments are routinely used in many settings, including the nuclear power industry, the chemical and petroleum industries, defense industries, the aerospace industry, food sciences, and health sciences. Industries that are increasingly using formal, quantitative methods of safety analysis include marine transportation and offshore systems, pipelines, motor vehicle, and recreational systems. The space program has stepped up its use of quantitative risk assessment since the *Challenger* accident. Other less publicized applications include the risk management program used by the U.S. Army for the disposal of chemical agents and munitions (Boyd and St. Pierre, 2001).

The government agencies most involved in using risk assessments are the U.S. Nuclear Regulatory Commission and the U.S. Environmental Protection Agency. Other agencies becoming active users of risk assessment methods are the U.S. Department of Energy, U.S. Department of Agriculture, U.S. Department of Defense, U.S. Food and Drug Administration, the National Aeronautics and Space Administration, and the U.S. Department of Transportation. The most active practitioners in the private sector are the nuclear, chemical, and petroleum

industries, although the scopes of application vary widely—the nuclear industry being the most consistent user of quantitative methods.

CONNECTING RISK ASSESSMENT AND DECISION ANALYSIS

Risk assessment is only one component of modern risk-management methods and decision analysis. A risk assessment is not a decision analysis, but advancements in risk assessment have contributed to quantitative decision analysis and provided a basis for more effective risk management. Risk assessment has also contributed to the development of analytical methods of quantifying other factors involved in decision making, such as costs and benefits. For example, the methods of treating uncertainties developed in quantitative risk assessment are also applicable to quantifying uncertainties associated with other factors. In the final analysis, the most important contribution of quantitative risk assessment to decision analysis is the quantification of low-probability, high-consequence events.

The foundation of QRA is the structuring of scenarios and methods of inferring the likelihood of events for which there is little or no actual experience. Bayes Theorem, a major advance in statistics, shows how to make better-informed decisions by mathematically blending new information with old information. This theorem has been fundamental to inferential thinking for both risk assessment and decision analysis (Bernstein, 1996).

Risk assessment and decision analysis have many buzzwords (e.g., Monte Carlo analyses, influence diagrams, multiple attributes, common-cause failures, realizations, minimum cut sets, sensitivity analyses, fault trees, event trees, etc.), but the basic principles are few. The principles focus on the development of scenarios describing how the system under study is supposed to work and scenarios indicating how the system can be made to fail, catastrophically or otherwise. The likelihood of events in the scenario must be linked to the supporting evidence. Events are propagated to an end state that terminates the scenario (i.e., the consequence). Other principles may be applied to aggregate the various end-states into the desired set of consequences.

The results of risk assessments are easy to interpret, including corrective actions having the biggest payoff in terms of risk reduction. Although the literature suggests many different risk assessment methodologies, in fact the differences are primarily in scope, application, boundary conditions, the degree of quantification, figures-of-merit, and quality. Like many other scientifically based methodologies, quantitative risk assessment is founded on relatively few basic principles.

REFERENCES

Apostolakis, G., R.J. Budnitz, P.O. Hedman, G.W. Parry, and R.W. Prugh. 1996. Report of the Risk Assessment Expert Panel on the Tooele Chemical Agent Disposal Facility Quantitative Risk Assessment. Prepared for Mitretek Systems. Falls Church, Va.

Apostolakis, G., H. Dezfuli, S.D. Gahring, M.B. Sattison, and W.E. Vesely. 2002. Report of the Independent Peer Review Panel on the Probabilistic Risk Assessment of the International Space Station: Phase II – Stage 7A Configuration. Washington, D.C.: National Aeronautics and Space Administration, Office of Safety and Mission Assurance.

Bernstein, P.L. 1996. Against the Gods: The Remarkable Story of Risk. New York: John Wiley & Sons.

Boyd, G.J., and G. St. Pierre. 2001. Risk Management Program for the Disposal of Chemical Agents and Munitions. Presented at Society for Risk Analysis: Special Symposium on Quantitative Risk Assessment. Sponsored by the Family Foundations of Chauncey Starr and B. John Garrick; B. John Garrick Foundation for the Advancement of the Risk Sciences, May 31-June 2, 2001. Laguna Beach, Ca.

Farmer, F.R. 1964. The Growth of Reactor Safety Criteria in the United Kingdom. Proceedings of the Anglo-Spanish Nuclear Power Symposium, Madrid, Spain.

Garrick, B.J. 1968. Unified Systems Safety Analysis for Nuclear Power Plants. Ph.D. thesis, University of California, Los Angeles.

Holmes and Narver, Inc. 1967. Reliability Analysis of Nuclear Power Plant Protective Systems. HN-190. Washington, D.C.: U.S. Atomic Energy Commission.

Jaynes, E.T. 2003. Probability Theory: The Logic of Science. Cambridge, U.K.: Cambridge University Press.

Jeffreys, H. 1957. Scientific Inference, 2nd Ed. Cambridge, U.K.: Cambridge University Press.

Kaplan, S., and B.J. Garrick. 1981. On the quantitative definition of risk. Risk Analysis 1(1): 11-27.

Kastenberg, W.E., G. E. Apostolakis, J.H. Bickel, R.M. Blond, S.J. Board, M. Epstein, P. Hofmann, F.K. King, S. Ostrach, J.W. Reed, R.L. Ritzman, J.W.Stetkar, T.G. Theofanous, R. Viskanta, and S. Guarro. 1988. Findings of the Peer Review Panel on the Draft Reactor Risk Reference Document. NUREG/CR-5113. Washington, D.C.: U.S. Nuclear Regulatory Commission.

PLG, Inc., Westinghouse Electric Corporation, and Fauske and Associates, Inc. 1981. Zion Probabilistic Safety Study. Prepared for Commonwealth Edison Company. Pittsburgh, Pa. Westinghouse Electric Corporation.

PLG, Inc., Westinghouse Electric Corporation, and Fauske and Associates, Inc. 1982. Indian Point Probabilistic Safety Study. Prepared for Consolidated Edison Company of New York, Inc., and the New York Power Authority. Pittsburgh, Pa.: Westinghouse Electric Corporation.

Raiffa, H. 1996. Decision Analysis. Columbus, Oh**:** McGraw-Hill Primis Custom Publishing.

Rechard, R.P. 1999. Historical relationship between performance assessment for radioactive waste disposal and other types of risk assessment. Risk Analysis 19(5): 763–807.

Starr, C. 1969. Societal benefits vs. technological risk. Science 168: 1232–1238.

USNRC (U.S. Nuclear Regulatory Commission). 1975. Reactor Safety Study: An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants. WASH-1400 (NUREG-75/014). Washington, D.C.: U.S. Atomic Energy Commission.

USNRC. 1990. Severe Accident Risks: An Assessment for Five U.S. Nuclear Power Plants: Final Summary Report. NUREG-1150 (Three Volumes). Washington, D.C.: U.S. Nuclear Regulatory Commission.

APPENDIX B

TESTIMONY OF PAUL H. GILBERT

Joint Hearing on
Implications of Power Blackouts for the Nation's Cybersecurity and Critical Infrastructure Protection: The Electrical Grid, Critical Interdependencies, Vulnerabilities and Readiness
Testimony of

Paul H. Gilbert, PE, NAE
Chairman
Panel on Energy Facilities, Cities, and Fixed Infrastructure
Committee on Science and Technology for Countering Terrorism
National Research Council
The National Academies
and
Director Emeritus of Parsons Brinckerhoff, Inc.

Before the

Cybersecurity, Science and Research and Development Subcommittee
and the
Infrastructure and Border Security Subcommittee
Select Committee on Homeland Security
U.S. House of Representatives

September 4, 2003

Good afternoon, Chairman Thornberry, Chairman Camp, and members of the Subcommittees. My name is Paul Gilbert. I am an officer and director emeritus of Parsons Brinckerhoff, Inc. I am also a member of the National Academy of Engineering and was Chair of the National Research Council Panel responsible for the Chapter on Energy Systems for the NRC Branscomb-Klausner Report: "Making the Nation Safer: the Role of Science and Technology in Countering Terrorism" <http://www.nap.edu/catalog/10415.html>. As you know, the NRC is the operating arm of the National Academy of Science, National Academy of Engineering and the Institute of Medicine, chartered in 1863, to advise the government on matters of science and technology. The subject report was the product of the mobilized academies following the 9/11 attacks. Some 130 volunteers from every branch of science, engineering and medicine assembled to undertake this work on an urgent basis with the report production financed entirely with private funds of the Academies. The report was first presented in June of 2002. It is a pleasure to

come before you today to assist in focusing attention on the vulnerabilities of our Electric Power Systems, including their cyber sub systems, and the enormous dependence of other critical infrastructure on the electric supply.

Our basic infrastructure systems are a highly integrated, mutually dependent generally highly utilized set of infrastructure components that provide our communities and way of life with vitally needed services and support. These include the electric power and our food supply, water supply, waste disposal, natural gas, communications, transportation, petroleum products, shelter, employment, medical support and emergency services, and all our other basic needs. While all these elements are essential to our well being, only one has the unique impact if lost of causing all the others to either be seriously degraded or completely lost. And that, of course, is electric power. Our technically advanced society is literally hard wired to a firm reliable electric supply.

That electric supply system has, over the past decade taken on significantly greater loads (power demands) and has also undergone a makeover from being a highly regulated, vertically integrated utility industry to one that is partially deregulated, far less unified and not so robust and resilient as it was. The generation side is essentially deregulated and operating under an open market set of conditions where competitive price, low operating costs and return on investment are rewarded with profits and bonuses. At the same time the transmission sector remains fully regulated and limited from taking steps to meet growing demand with new capacity by uncertainty in knowing how such investments will be paid for under regulatory bodies that are tasked to see that power is delivered to rate payers at minimum cost. Where possible, operating costs have been reduced by installing automated cyber controllers, SCADA units and LANs, to perform the functions that people had previously performed. In general, control is now more centralized, spare parts inventories are reduced, and systems are highly integrated across entire regions.

This dramatic change has played out with the result that the in-place electrical systems assets today are typically being operated very efficiently at close to the limit of available capacity. In this mode, another characteristic of such systems appears. When operated near their capacity, these systems have little margin within which to handle power or load fluctuations. Thus they are quite vulnerable to being brought down by operating fluctuations that exceed their remaining margins. Shutting down becomes the only way a system element has of protecting itself from severe damage when load exceeds capacity. But the loss of a piece of the grid, a section of transmission line, does not end the problem. The line down takes with it the power it was transmitting. A connected power plant, having no connected load must also shut down. In these highly integrated grids, more lines have imbalance problems and more plants sense capacity problems and so also shut down. This cascading spreads very rapidly in many directions and in seconds, an entire sector of the North American grid can be down. We had a living example of this event, this past month, caused by an accident. We were fortunate to see the power return in so short a time.

The exact same consequences could too easily be reproduced by an attack from a small trained terrorist team as was hypothecated in the Making the Nation Safer report. Several critical nodes in the grid, taken out in the most damaging manner is the terrorist attack. What is caused is

the terror flowing to all of us from the attack. Recovery in the case cited might take weeks or months, not hours or days, and the damage done to our people and our economy would be enormous.

While the report does not speculate on the extended consequences of such an event, I have been asked to do so here and so offer this as personal opinion. Because our critical infrastructure is so completely integrated, with the power out for even a day or two, both food and water supply soon fail. Transportation systems would be at a standstill. Wastewater could not be pumped away and so would become a health problem. In time natural gas pressure would decline and some would loose gas altogether. Nights would be very dark and communications would be spotty or non-existent. Storage batteries would have been long gone from the stores if any stores were open. Work, jobs, employment, business and production would be stopped. Our economy would take a major hit. All in all our cities would not be very nice places to be. Some local power grids would get back up and so there would be islands of light in the darkness. Haves and have-nots would get involved. It would not be a very safe place to be either. Marshal Law would likely follow along with emergency food and water supply relief. We would rally and find ways to get by while the system is being repaired. In time, the power will start to come back. Tentatively at first, with rolling blackouts and then with all its glory. Several weeks to months have passed, and the clean up would begin. This is one man's opinion.

We have the means to limit the kind of disaster that has been speculated upon above. The recommendations provided in Chapter 6 of the report address actions that are designed to minimize the immediate vulnerabilities of the electric power systems and then to seek longer-term solutions. Those recommendations are as to the point today as they were when published 15 months ago.

- The recommendations begin with immediate attention is needed to mobilize the leadership and then the resources of people and organizations to first determine the proper roles for each interested party and then to come together, meet and develop needed plans.
- Issues that deter open discussions among the private and governmental parties need to be resolved immediately. These include antitrust, liability and FOIA.
- Review by government of the institutional and market settings (regulated and deregulated and open free market) for the industry needs attention to focus the inherent incentives on what the nation needs to live safely.
- Mobilization of tools now employed by the military to analyze vulnerabilities should occur, perhaps transferring them to DHS for use with the grids.
- Coordinated studies are indicated to identify the most critical equipment in the respective power systems and to describe the protective measures to be taken with each.
- Simulation models of these highly complex grids are indicated that are capable of identifying points of greatest vulnerability and reserves on operating capacities.
- Statutory action is indicated to allow recovery crews to immediately enter what would then be a crime scene following an attack to commence the work of repair, recovery, and restoration of service.

- The regulatory bodies must be encouraged to find the means for transmission organizations to define costs for counter terrorism improvements and for recovering those costs from their operations or from other sources.
- The use of SCADA systems in an unprotected configuration should be addressed and expert advice obtained regarding the options available to correct the vulnerabilities now present.
- Research is indicated that addresses particular system equipment needs. First among the list is the potential value of modular universal EHV transformers to support rapid grid recovery.
- Research is indicated into the equipment and technology required for, and the steps involved to, transition to an intelligent, adaptive power grid.

There is much greater detail and substance provided in Chapter 6 of the referenced report. The unfortunate black out this past month has drawn important attention to this area of critical infrastructure need. We at the Academies are delighted that we can continue to contribute to the effective resolution of these issues.

Thank you for inviting me today and for your attention in holding these hearings. I will be happy to respond to your questions.

APPENDIX C

STUDY GROUP BIOGRAPHIES

B. John Garrick, ***chair,*** independent consultant, was a cofounder of PLG, Inc., an international engineering, applied science, and management consulting firm, from which he retired as president and chief executive officer in 1997. His professional interests include risk assessment in nuclear energy, space and defense, chemicals and petroleum, and transportation. A past president of the Society for Risk Analysis, Dr. Garrick is also a fellow of three professional societies and a member of the National Academy of Engineering. He has received numerous awards, including the Society for Risk Analysis Distinguished Achievement Award. Dr. Garrick was appointed to the U.S. Nuclear Regulatory Commission Advisory Committee on Nuclear Waste in 1994 and is the current chair. He has served on several National Research Council Committees and was chair of the Committee on the Waste Isolation Pilot Plant. Dr. Garrick received his B.S. in physics from Brigham Young University and his M.S. and Ph.D. in engineering and applied science from the University of California, Los Angeles; he is also a graduate of the Oak Ridge School of Reactor Technology.

James E. Hall, chair of the National Transportation Safety Board (NTSB) from 1994 to 2001, has worked tirelessly throughout his career to improve the safety of all modes of transportation. During his term of office at NTSB, he was chairman of the Board of Inquiry for many major accidents, including the crashes of USAir 427 and TWA 800; Mr. Hall also represented the NTSB internationally in numerous investigations. During his chairmanship, the NTSB issued landmark safety studies on commuter airlines, the air tourist industry, the performance and use of child restraint systems, the dangers to children of passenger-side air bags, personal watercraft safety, oversight of transit buses, and the safety of passive grade crossings. In September 1996, President Clinton named Mr. Hall to the White House Commission on Aviation Safety and Security, which issued two reports recommending improvements in aviation safety and security around the world. Mr. Hall received a law degree from the University of Tennessee and was awarded a Bronze Star for Meritorious Service for his service in Vietnam. He is currently affiliated with Hall and Associates based in Washington, D.C., and Chattanooga, Tennessee.

Max Kilger is the social psychologist on the Honeynet Project, a 30-member international team conducting primary research on computer security. Dr. Kilger's work is focused on the behavior and motivations of individuals involved in cyberattacks on networked computer systems. He has extensive experience in quantitative and qualitative research methodologies and is involved in developing techniques for integrating disparate databases for the purpose of predicting behavior in scarce information environments. Dr. Kilger earned his Ph.D. from Stanford University in social psychology. During his teaching career at San Jose State University and City University of New York, he taught courses on statistics and research methodology and developed a popular course on relationships between people and technology. He is currently director of Statistical Sciences for Symmetrical Resources, a national research company.

John C. McDonald is founder and president of MBX, Inc., a communications research organization. In 1991, he was chief scientist for Contel Corporation, which he joined in 1970;

his work there included managing development teams in local message metering and digital switching systems, all based on his inventions. Mr. McDonald's work at Contel led to the pioneering installation of the first public digital switched-telecommunications network in North America. Early in his career, he worked for GTE Sylvania, where he was co-inventor of Doppler radar systems. Mr. McDonald has been a member of the advisory boards of Stanford University, Manhattan College, Polytechnic University, and the University of Maryland. He is a member of the National Academy of Engineering (NAE) and has served on the National Research Council (NRC) Board on Communications and Computer Applications and many NRC committees. He is a member of Tau Beta Pi, Sigma Xi, and past president of the IEEE Communications Society, a life fellow of the IEEE, and a recipient of the IEEE Donald W. McLellan Award and Centennial Medal. He is a registered professional engineer and a credentialed California teacher. Mr. McDonald earned his B.S, M.S., and D.E.E. in electrical engineering from Stanford University. He has published more than 80 technical papers, two books, and an encyclopedia of telecommunications. He also holds 20 patents.

Tara O'Toole is CEO and director of the Center for Biosecurity at the University of Pittsburgh Medical Center. She previously served as director of Johns Hopkins Center for Civilian Biodefense and was a faculty member at Johns Hopkins Bloomberg School of Public Health in the Department of Health Policy and Management. Dr. O'Toole serves in many other advisory and consulting positions related to bioterrorism preparedness. She is a member of the Defense Science Board Summer Panel on biodefense technologies and the Maryland Department of Health and Mental Hygiene steering group on public health response to weapons of mass destruction and co-editor in chief of a new journal, *Biosecurity and Bioterrorism – Biodefense Strategy, Practice and Science.* From 1993 to 1999, Dr. O'Toole was assistant secretary of energy for environmental safety and health. She is board-certified in internal medicine and occupational medicine. She received a B.A. from Vassar College, an M.D. from George Washington University, and an M.P.H. from Johns Hopkins University.

Peter S. Probst worked for the Central Intelligence Agency (CIA) and the Office of the Secretary of Defense for almost 30 years. He is currently a private consultant and vice president and director of programs for the Institute for the Study of Terrorism and Political Violence and coauthor of *Terror-2000: The Future Face of Terrorism*. Mr. Probst is a member of the American Society of Industrial Security Standing Committee on Global Terrorism, Political Instability and International Crime, a member of the Advisory Board of the Investigative Project on Religious Extremism sponsored by the Middle East Forum, and a member of the International Research Group on Terrorism sponsored jointly by the British Airey-Neave Trust and the U.S. Institute for Peace. He has been an invited speaker at private institutions, civic organizations, and governments in the United States and many other countries. Mr. Probst received a B.A. in history from Columbia College and an M.A. in anthropology/archaeology from Columbia University.

Elizabeth Rindskopf Parker is currently dean of the McGeorge School of Law, University of the Pacific. Prior to this, she was general counsel for the University of Wisconsin System and counsel to the international law firm of Bryan Cave, LLP. Her expertise is in public policy and international trade issues, particularly technology transfer. Her earlier experience includes general counsel for the Central Intelligence Agency; principal deputy legal adviser, U.S.

Department of State; general counsel, National Security Agency, U.S. Department of Defense; and acting assistant director (mergers and acquisitions), Federal Trade Commission. A graduate of the University of Michigan, Dr. Rindskopf is a frequent speaker on law and national security and has been a visiting professor at Case Western Reserve Law School and Cleveland State School of Law.

Robert Rosenthal recently joined Booz Allen Hamilton after retiring from the National Institute of Standards and Technology, where he had a 30-year career in research on computer networks and security, standards for public infrastructures, security communication protocols, high-confidence systems and software, critical infrastructure protection, and local and large-scale networking. During the Clinton administration, he worked closely with the National Security Council and the President's Commission on Critical Infrastructure Protection; he also served on numerous White House committees and working groups, including the Office of Science and Technology Policy's working groups on high-confidence systems and software, large-scale networking, and critical infrastructure protection. He was a program manager for portions of the survivable large-scale systems and high-confidence networking programs at the Defense Advanced Research Projects Agency and program manager for the Computer Emergency Response Team at the Software Engineering Institute, Carnegie Mellon University. Dr. Rosenthal is a past chairman of the IEEE Technical Committee on Computer Communications and a past member of Sigma Xi.

Alvin W. Trivelpiece, now director emeritus, Oak Ridge National Laboratory (ORNL), was laboratory director from January 1989 through March 2000 and, also in 1989, vice president of Martin Marietta Corporation, the managing and operating contractor for ORNL. In January 1996, he was appointed president of Lockheed Martin Energy Research Corporation, the new managing and operating contractor for ORNL. From April 1987 to January 1989, Dr. Trivelpiece was the executive officer of the American Association for the Advancement of Science (AAAS), the country's leading general science organization. Prior to taking on his responsibilities at AAAS, he was the director of the Office of Energy Research at the U.S. Department of Energy, corporate vice president at Science Applications International, Inc., and vice president for engineering and research at Maxwell Laboratories. Dr. Trivelpiece's academic experience includes: professor of physics at the University of Maryland and professor, Department of Electrical Engineering, University of California, Berkeley. While on leave from the University of Maryland, he was assistant director for research in the Division of Controlled Thermonuclear Research, U.S. Atomic Energy Commission. He received his B.S. from California Polytechnic State University and his M.S. and Ph.D. from the California Institute of Technology. He holds several patents on accelerators and microwave devices and is the author or coauthor of many technical reports and two books. Dr. Trivelpiece was elected to the National Academy of Engineering in 1993.

Lee A. Van Arsdale is currently president of Unconventional Solutions, Inc., a private consulting firm that focuses on terrorism and security matters. He graduated from the U.S. Military Academy at West Point in 1974 and served in the U.S. Army until his retirement in 1999. In the course of his Army career, Mr. Van Arsdale served in three combat zones in command positions and was decorated for valor with the Silver Star and the Purple Heart for wounds received in combat. He also participated in numerous classified operations around the

world. During his final two years in uniform, he served at the Pentagon as the Counterterrorism/Special Projects Branch Chief in the Office of the Assistant Secretary of Defense for Special Operations and Low-Intensity Conflict. He received a B.S. in engineering from West Point and an M.S. from the University of Colorado, and he is a graduate of the Armed Forces Staff College and the U.S. Army War College (AWC), where he won the AWC Foundation Excellence in Writing Award. Mr. Van Arsdale is a member of the Special Operations Association and Mensa.

Edwin L. Zebroski, a consultant on risk analysis and decision analysis (through Elgis Consulting Company), has extensive experience in the design, development, safety, materials, fuel cycle, and economic aspects of power systems. He is the author of more than 150 technical publications, mostly on energy-related topics, including patents and sections of seven books. His previous positions include manager of development engineering, General Electric Company; director of the Systems and Materials Department and chief nuclear scientist at the Electric Power Research Institute; vice president of engineering at the Institute for Nuclear Power Operation; and director of Risk Management Services at Aptech Engineering Company. He received a B.S. in physics and chemistry from the University of Chicago and a Ph.D. in physical chemistry from the University of California, Berkeley. Dr. Zebroski was elected to the National Academy of Engineering in 1981.

APPENDIX D

BRIEFINGS

The following individuals briefed one or more study group members.

Massoud Amin
Area Manager, Infrastructure Security, Grid Ops/Planning and Markets/Lead, Math and Information Sciences
Electrical Power Research Institute

Vulnerability of the power grid system to terrorist attack

Douglas Bauer
Director for Counterterrorism at the National Academies
National Research Council

Review of the National Academies counterterrorism initiatives

Louis Branscomb
Professor Emeritus, Public Policy and Corporate Management
John F. Kennedy School of Government
Harvard University

National Research Council Report: Making the Nation Safer

Todd Brethauer
Science Advisor, Technical Support Working Group
WINTEC SETA

Federal agency perspectives on combating terrorism

John Carlson
Senior Advisor and Acting Director
Office of the Comptroller of the Treasury

An intelligence user agency perspective on terrorist threats

Larry Fogel
President
Natural Selections

Applying risk analysis to counterterrorism

B. John Garrick
Independent Consultant

Overview of risk assessment methodologies

Sandy Hauserman
Vice President
Guy Carpenter

Risk assessment: perspective of the re-insurance industry

Stanley Kaplan
Partner
Bayesian Systems

Quantitative risk assessment

Donald M. Kerr
Deputy Director for Science and Technology
Central Intelligence Agency

Federal agency perspectives on combating terrorism

Richard Klausner
Executive Director, Senior Fellow, and Special Advisor to the President of the National Academies for Counter Terrorism and Co-chair, Committee on Science and Technology for Countering Terrorism (until 5/31/02)
Executive Director, Global Health Program, Bill and Melinda Gates Foundation

Review of the National Academies counterterrorism initiatives

John Knight
Professor of Computer Science
University of Virginia

Terrorism and critical infrastructure sector vulnerability: a software engineering perspective

Rennselaer Lee
Senior Terrorism Policy Consultant
Congressional Research Service

Keynote address: The changing nature of terrorism

Ronald Lehman
Director of the Center for Global Security Research
Lawrence Livermore National Laboratory

Federal agency perspective on combating terrorism

John Marburger
Director of the Office of Science and Technology Policy
Office of the President

Science, engineering, technology, and counterterrorism

Mike McConnell
Director, Infrastructure Assurance Center
Booz Allen Hamilton

Ensuring infrastructure security: vulnerability assessment and policy issues

Ltg. Duncan McGill
Chief Consequence Management Branch, Emergency Management Division

Defense Threat Reduction Agency

Federal agency perspective on combating terrorism

Richard Meserve
Chairman
U.S. Nuclear Regulatory Commission

Nuclear plant preparedness and security

David L. Osbourne
Head, Research Section
Federal Research Division, Library of Congress

Psychological Motivations of Terrorists

Suzanne Spaulding
Chair
American Bar Association Standing Committee on Law and National Security

Terrorism, emerging policy issues, trends, future directions, societal strengths and weaknesses, recommendations of commissions, unresolved issues

John H. Stevens
Manager, Project World Tec
Defense Intelligence Agency

Application of Risk Analysis to Power Grid Vulnerabilities

Ronald Taylor
Study Director, Committee on Science and Technology for Countering Terrorism
Department on Engineering and Physical Systems
National Research Council

Review of the National Academies counterterrorism initiatives

Richard Wilson
Mallinckrodt Professor of Physics
Harvard University

Role of risk assessment in public policy

Lee M. Zeichner
President
Risk Analytics, LLC

Critical infrastructure sector vulnerability: a software engineering systems overview