RISK ASSESSMENT OF EXTREME EVENTS

Rae Zimmerman
Professor of Planning and Public Administration and Director, Institute for Civil Infrastructure Systems (ICIS) at New York University's Robert F. Wagner Graduate School of Public Service
E-mail: Rae.zimmerman@nyu.edu.

Vicki M. Bier
Professor in the Departments of Industrial Engineering and Engineering Physics and Director of the Center for Human Performance and Risk Analysis at the University of Wisconsin-Madison
E-mail: bier@engr.wisc.edu.

I. Introduction: Scope

Risk assessment is a means to characterize and reduce uncertainty to support our ability to deal with catastrophe through risk management. This paper addresses the application of risk assessment to both the built and natural environments, to improve our understanding and management of human health, safety, and security within those environments, especially in the context of extreme and unanticipated events.

Modern risk assessment began over three decades ago with applications in the military and nuclear power, beginning with the Reactor Safety Study (U.S. Nuclear Regulatory Commission, 1975). In the late 1970s it gradually expanded, and was applied to the vast array of chemical risks being regulated under dozens of federal environmental statutes. For example, risk assessment has been used as the foundation for: setting drinking water, ambient water quality, and air quality standards; review and renewal of pesticide applications; and determining levels of site cleanup under the Superfund program. Applications to engineered systems, and in particular infrastructure, are common; examples are given in Moses (1998) and Lave and Balvanyos (1998). Blockley (1992) also devotes a number of chapters to civil engineering topics (e.g., design codes or risk assessment in structural engineering), and several infrastructure-engineering applications (e.g., dam safety, marine structures).

II. What is Risk Assessment?

Risk assessment connotes "a systematic approach to organizing and analyzing scientific knowledge and information for potentially hazardous activities or for substances that might pose risks under specified circumstances" (National Research Council 1994: 4). This definition reflects the flexibility that has been incorporated into the concept over the years since it was first introduced.

The overall objective of risk assessment is to estimate the level of risk associated with adverse health and safety effects of stressors or sources of those effects. Health and safety effects at the most general level encompass both human and ecological health. Stressors are the broad set of causes, such as chemicals, biological agents, and physical conditions in the built environment and the physical structures that pervade our everyday life. Between cause and effect are numerous steps that link the two, and risk assessment identifies and quantifies those steps.

Although the steps vary by the setting, the sources, the effects, and the concepts are similar at a generic level.

A. Definitions of Risk

In their seminal paper, Kaplan and Garrick (1981) define "risk" as involving "both uncertainty and some kind of loss or damage." More specifically, Kaplan and Garrick note that "Risk includes the likelihood of conversion of that source into actual…loss, injury, or some form of damage." Thus, fundamentally, risk depends both on the probability or frequency of an adverse outcome, and also on the severity of that outcome. Risk has similarly been defined generally as "the potential for realization of unwanted, negative consequences of an event" (Rowe 1976: 24). More quantitatively, Sage and White (1980: 426) define risk as "the probability per unit time of the occurrence of a unit cost burden," and state that it "represents the statistical likelihood of a randomly exposed individual being adversely affected by some hazardous event." Thus, risk has been defined at many different levels of detail.

More debate has occurred over acceptable levels of risk than over the definition itself. The debate has occurred with respect to the process of deriving, selecting, and applying acceptable levels of risk, including determination of who makes the decision. As a consequence of those debates not being addressed up front, the passage of environmental health risk statutes over the past two decades has spawned a large number of terms to characterize acceptable risk levels. These terms include: "adequate"; "imminent"; "substantial"; "reasonable" vs. "unreasonable"; "posing grave danger"; "at a zero level"; "significant" vs. "de minimus"; and "an ample or adequate margin of safety." More important than the differences in terminology is that different statutes can allow different levels of risk for the same risk factor or agent, such as a hazardous chemical.

B. Relationship of Risk to Other Concepts

The concept of risk is related to a number of other concepts, including threat, hazard, vulnerability, uncertainty, and variability. In addition, terms related to the degree of risk (such as "catastrophes," "extreme events," "unexpected events," and "unanticipated events") are also common. It is important to sort these out to ensure a full understanding of how they relate to risk assessment.

A "hazard" has been defined as "An act or phenomenon that has the potential to produce harm or other undesirable consequences to humans or what they value" (Stern and Fineberg 1996: 215). Similarly, Merriam-Webster's Collegiate Dictionary (2002) defines "hazard" simply as "a source of danger." Kaplan and Garrick illustrate the distinction between hazard and risk by considering the ocean as a hazard, and noting that the degree of risk undertaken in crossing the ocean will depend on the level of safeguards adopted (e.g., the method of transportation used).

"Threat" is another concept closely linked to risk. The definition of threat is often expressed in the context of vulnerability. The National Research Council (2002) defines "vulnerability" as "an error or a weakness in the design, implementation, or operation of a system." The National Research Council goes on to define "threat" as "an adversary that is motivated to exploit a system vulnerability and capable of doing so," and "risk" as "the likelihood that a vulnerability will be exploited, or that a threat may become harmful." Thus, their definition of risk is limited to intentional threats, unlike the more general definition given above.

Closely related to the idea of risk are the ideas of extreme, catastrophic, and/or unexpected events. These are critical in linking risk assessment to risk management, since they are value-laden terms that reflect human perceptions. Bier et al. (1999) define "extreme events" as being extreme in terms of both their low frequency and high severity. They note that an event may be rare but not extreme: "For example, a room temperature of exactly 72.34571° F may occur extremely rarely, but would hardly seem to be extreme by most people's definitions." Similarly, a "catastrophe" is defined (Merriam-Webster Collegiate Dictionary, 2002) as "a momentous tragic event ranging from extreme misfortune to utter overthrow or ruin." Catastrophic risks are often distinguished from "chronic" risks, defined by Merriam-Webster as "marked by long duration or frequent recurrence." Thus, for example, asbestosis would generally be considered a chronic risk, since it develops gradually over a long period of time, and then persists. Finally, Shackle (1972) distinguishes between "counterexpected events" (events that had been identified, but were believed to be unlikely or impossible) and "unexpected events" (which had not even been anticipated). Unexpected events are obviously a significant challenge for both risk analysis and risk management.

It is important to clarify the distinction between the concepts of "uncertainty" and "variability." In particular, Kaplan (1983) distinguishes between "state of knowledge" uncertainty and "population variability." To illustrate, in the context of equipment failure rates, state of knowledge uncertainty may refer to a lack of knowledge about the average or mean failure rate across the entire population of similar items of equipment. By contrast, population variability would refer to differences among the items in the population. Similarly, in health risk assessment, state of knowledge uncertainty may refer to a lack of knowledge about the average effect of a particular chemical on health, while population variability would refer to differences in susceptibility from one person to another. Thus, one can have state of knowledge uncertainty even without population variability (e.g., if all items or people in a population have the same unknown risk level), and also population variability without state of knowledge uncertainty (e.g., if the risks to different individuals in the population are known but unequal). This distinction is discussed at length in Bier (2001), Cullen and Frey (1999), Morgan and Henrion (1990), and the National Research Council (1994). Similarly, in the context of health risk assessment, state of knowledge uncertainty may refer to lack of knowledge about the average effect of a particular chemical on human health, while population variability refers to differences in susceptibility to the chemical from one person to the next.

C. Systems Analysis Paradigms and Frameworks for Risk Assessment

In their initial formulations, risk assessment and risk management were distinctly separate functions. The so-called "Red Book" (National Research Council 1983: 7) clearly distinguishes risk assessment and management in the context of risk regulation, recommending that:

> regulatory agencies take steps to establish and maintain a clear conceptual distinction between assessment of risks and consideration of risk management

> alternatives; that is, the scientific findings and policy judgments embodied in risk assessments should be explicitly distinguished from the political, economic, and technical considerations that influence the design and choice of regulatory strategies.

The report, however, clearly recommended that the two processes not be placed in separate agencies.

Over time, however, fields such as environmental sociology and social psychology infused issues of risk perception and risk communication into risk assessment, profoundly influencing the nature of risk assessment. In the 1990s, the two processes began to be more formally merged (this section is drawn from Zimmerman 1998). The National Research Council (NRC) Committee on Risk Assessment of Hazardous Air Pollutants *Science and Judgment* report in 1994 (National Research Council, 1994: 24), the NRC *Understanding Risk* study (Stern and Fineberg 1996: 28), the U.S. Environmental Protection Agency (EPA) strategic plan for its Office of Research and Development (U.S. EPA May 1996: 3), and the report of the Presidential/ Congressional Commission on Risk Assessment and Risk Management (1997, Volume 2: 7) each reflect far more interaction between stakeholders and risk assessors than the approach described in the "Red Book." The manner in which the various risk assessment steps are integrated with stakeholder and risk management concerns varies, however, among these reports. Stern and Fineberg (1996) develop this integration still further, with the introduction of concepts of analysis and deliberation, both of which are seen as part of the fully integrated risk assessment/risk management paradigm, and viewed as iterative processes:

Analysis, according to Stern and Fineberg (1996: 214) is defined as "The systematic application of specific theories and methods, including those from natural science, social science, engineering, decision science, logic, mathematics, and law, for the purpose of collecting and interpreting data and drawing conclusions about phenomena," and "uses rigorous, replicable methods, evaluated under the agreed protocols of an expert community-such as those of disciplines in the natural, social, or decision sciences, as well as mathematics, logic, and law- to arrive at answers to factual questions" (Stern and Fineberg 1996: 4). Conversely, they state: "Deliberation uses processes such as discussion, reflection, and persuasion to communicate, raise and collectively consider issues, increase understanding, and arrive at substantive decisions" (Stern and Fineberg 1996: 20). The combination or linkage of these two steps is termed the "analytic-deliberative" process, an iterative process in which deliberation and analysis are viewed as complementary (Stern and Fineberg 1996: 3), so that "deliberation frames analysis and analysis informs deliberation" (Stern and Fineberg 1996: 20).

To understand how the integration of these processes and the changes in the boundaries between them have occurred, it is instructive to view risk assessment and risk management from a systems perspective. From a systems perspective, the boundaries between risk assessment and risk management are seen to be flexible. For example, interested stakeholders under such a framework are more fully included in the process, and can provide inputs into the design of a risk assessment at the very outset; e.g., in the selection of exposure points and health outcomes. In the aftermath of September 11th, this in fact has been occurring with respect to air quality and its effect on human health. Residents and workers have voiced strong concern over the extent to which the design of the air quality sampling strategy (including the placement of monitors and the selection of standards used to interpret the data) reflect what they are experiencing. The events following the destruction of the buildings and their impact on air quality have little precedent to guide the controversy.

Fundamentally, risk assessment is a form of systems analysis, and thus follows the systems approach. Haimes (1998; see also Haimes and Schneiter 1996) lays out the steps in the systems approach **(**including the involvement of stakeholders**)**, as paraphrased below:

1. Identify the problem
2. Determine the objectives and goals of the analysis
3. Become familiar with the total problem environment
4. Study the system being analyzed
5. Develop models (including multiple models, if necessary)
6. Solve the models
7. Identify various feasible solutions to the identified problem
8. Evaluate the costs and benefits of the identified solutions
9. Communicate the solutions to the client
10. Evaluate the impact of current decisions
11. Implement the selected solution
12. Review the analysis
13. Iterate

Kaplan and Garrick (1981) state that the goal of a risk assessment is to answer three questions: "What can go wrong?" "How likely is it that that will happen?" "If it does happen, what are the consequences?" Thus, strictly speaking, their definition of risk goes through step 7 in the above list, with the remaining steps falling into the category of "risk management." However, most risk assessments in practice do include the identification and evaluation of possible risk reduction measures; that is, risk assessment and risk management are integrated to a much greater extent than the portrayal of the process as independent steps would lead one to believe.

The general framework presented above is applicable to both health/environmental and engineering risk assessment. However, some of the details of the methodology differ between the two areas. Therefore, the next section presents a discussion of risk assessment methodologies for each of these areas separately.

III. Risk Assessment Methodologies for Health and Engineered/Built Systems

A. Health and Environmental Risk Assessment

The steps in risk assessment formalized in 1983 in the "Red Book" include hazard identification, dose-response assessment, exposure assessment, and risk characterization. The latter three steps are sometimes referred to as risk estimation, to distinguish them from the more qualitative process of hazard assessment. Although these generic steps have for the most part remained unchanged since the early 1980s, and have been applied in some of the major environmental statutes (National Research Council 1994), their interpretation and integration with risk management has undergone dramatic change. This is reflected in the change in the term used in the 1996 *Understanding Risk* report (Stern and Fineberg 1996), from “risk assessment” to “risk analysis,” to address the breakdown between the production of scientific results and their use by decision makers and managers (Stern and Fineberg 1996: x). Assessment is usually used in a way that implies quantification of a rather narrow set of things, whereas analysis is a broader concept that connotes a separation of the whole into its components to obtain new knowledge. The components of risk assessment (as used in health and environmental risk assessment) are described below.

Hazard Identification

The objective of hazard identification in the health and environmental area is to determine which hazards are to be used as the subject of the risk estimation. Hazard identification in the health and environmental area draws upon methods such as epidemiology (to identify disease clusters), structure-activity relationships (linking the structure of one chemical to the possible or likely biological effects of another similar chemical), and short-term tests.

Risk Estimation

A generic equation for the calculation of the risk estimate is as follows:

Risk = Unit Risk Level (derived from dose-response studies) x Exposure Level

Exposure Assessment

Exposure assessment is used to calculate the level of a hazardous substance to which people (and/or the environment) may be exposed, and is part of consequence assessment. Exposure assessment typically is comprised of two distinct stages—one having to do with the migration and fate of contaminants from sources to sinks, and another having to do with the exposures of individuals. Uncertainty exists at every stage of the assessment, especially where default values are used (U.S. EPA, 1999); however, few risk assessments quantify the assessment at each stage to derive a composite uncertainty estimate. A generic equation for estimating exposure levels is as follows, similar to the ones used in Superfund risk assessments (U.S. EPA, 1989):

$$\text{Exposure Level} = \frac{\text{Contaminant Concentration x Contact Rate x Exposure Duration}}{\text{Body Weight x Lifetime}}$$

Dose-Response

Dose-response relationships are the foundation for estimating a risk level for a given organism and risk agent under specific conditions. For health and environmental data, this is usually expressed in terms of a dose-response curve; as one increases the dose, the response changes accordingly. In this step—the construction of dose-response relationships—many uncertainties appear. For example, the shape of such a curve is easier to derive at higher doses, but at very low doses, the shape of the curve has to be modeled which introduces considerable uncertainty.

Risk Characterization

Stern and Fineberg (1996: 216) define risk characterization as "A synthesis and summary of information about a hazard that addresses the needs and interests of decision makers and of interested and affected parties. Risk characterization is a prelude to decision making and depends on an iterative, analytic-deliberative process."

B. Engineered Systems

Like health and environmental risk assessment, the process of risk assessment for engineered systems begins with hazard identification. In the context of engineering risk assessment, this generally includes system familiarization, and sometimes also qualitative methods such as hazard and operability studies or failure modes and effects analysis. Hazard identification is followed by the assessment of accident occurrence frequencies, and consequence analysis (including the use of dose-response models as appropriate). As noted above with regard to health risk assessment, uncertainty analysis is an overarching step that encompasses the entire process, since uncertainties can arise with regard to both accident frequencies and consequences.

Assessment of Accident Occurrence Frequencies

Bier (1997) summarizes the process of probabilistic risk assessment (PRA) for complex engineered systems. To structure the list of possible scenarios for what can go wrong, PRA models are generally hierarchical in nature. Two reliability analysis techniques are commonly used to represent the combinations of component failures that could lead to an accident: fault trees; and event trees (McCormick 1981). (Reliability block diagrams, which are logically equivalent to fault trees, are sometimes also used.) Fault trees and event trees are equivalent in the sense that it is possible to represent the same system or subsystem either way, and both techniques are useful. However, each technique has different strengths and weaknesses (Pate-Cornell 1984), and they therefore tend to be used in different contexts.

As noted by Pate-Cornell, fault trees are constructed using inductive or "backward" logic. In other words, the process starts with a hypothesized system or subsystem failure, and works backwards to identify which combinations of component failures could give rise to that event. By contrast, event trees are constructed using deductive or "forward" logic. In other words, rather than hypothesizing a system failure, the process starts by hypothesizing an initiating event (i.e., a departure from normal operations), then identifies all possible combinations of subsequent events (i.e., successes or failures of particular components or subsystems), and determines which sequences of events could cause failure of the system as a whole. Figure 1 shows a simplified event tree representing an initiating event and the subsequent response of four subsystems (A, B, C, and D). For each subsystem, the upper branch represents success and the lower branch represents failure. Thus, the event sequence shown in bold consists of the initiating event I, followed by success of subsystem A, failure of subsystem B, success of subsystem C, and failure of subsystem D.

The frequency of this scenario, S, can be quantified as a product of terms according to

$$F(S) = F(I)\,P(A|I)\,P(B'|IA)\,P(C|IAB')\,P(D'|IAB'C)$$

where:

$F(S)$	=	Frequency of scenario S;
$F(I)$	=	Frequency of initiating event I;
$P(A\|I)$	=	Conditional probability that subsystem A succeeds given that initiating event I has happened;
$P(B'\|IA)$	=	Conditional probability that subsystem B fails given that initiating event I has happened and subsystem A has succeeded;
$P(C\|IAB')$	=	Conditional probability that subsystem C succeeds given that initiating event I has happened, subsystem A has succeeded, and subsystem B has failed; and
$P(D'\|IAB'C)$	=	Conditional probability that subsystem D fails given that initiating event I has happened, subsystem A has succeeded, subsystem B has failed, and subsystem C has succeeded.

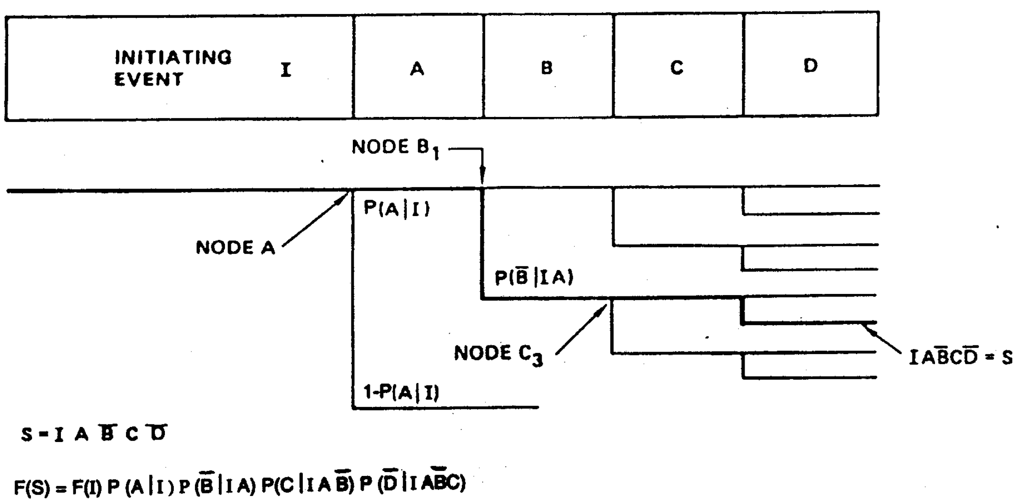

Figure 1. Sample Event Tree

The conditional probabilities P(•) represent the long-run fraction of times that a scenario will follow a particular branch of the event tree. The conditional nature of these split fractions reflects the fact that subsystems may be dependent on each other. For example, if subsystem A provides electric power to subsystem B, then the failure probability of subsystem B will depend on whether subsystem A has succeeded or failed. This is also true for the initiating event I; for example, if the initiating event is an earthquake, this may affect the failure probabilities of other subsystems.

As mentioned above, while both fault trees and event trees are useful, they have different strengths and weaknesses. In particular, event trees are well suited for displaying the order of events, and also dependencies between events (e.g., the fact that the failure probability of subsystem B may depend on the status of subsystem A). Event trees are therefore useful for facilitating communication about the assumptions made in the risk model; e.g., for presenting a risk model to plant staff for review and discussion. However, because combinations of subsystem successes and failures are shown explicitly, event tree models can rapidly become extremely large, including literally billions of sequences. Fault trees, by contrast, provide a more compact way of representing combinations of events, but can obscure dependencies and the chronological order in which events unfold. Therefore, most PRA's use event trees to structure the overall risk assessment. Fault trees can then be used to quantify the failure probabilities of particular systems or subsystems represented in the event tree.

Available data is then used to help quantify the various quantities that appear in the risk model. These include: initiating event frequencies; component failure rates (including the likelihood of common cause failure, in which multiple components fail during a short period of time for the same reason); maintenance frequencies and durations (to reflect the fact that components are temporarily out of service while they are being maintained); and human error rates. The data used to support the quantification process can include component-specific information for the specific items of equipment being considered, generic information (such as failure rates for similar components), and expert opinion. Success or exposure data is generally needed in addition to failure data, to provide the denominators for the desired failure rates.

Bayesian statistics (Martz and Waller 1982) is frequently used for data analysis in support of risk assessments, since Bayesian statistics are well suited to the analysis of sparse data, and provide probability distributions for the quantities of interest in addition to point estimates (thus supporting the overarching task of uncertainty analysis). Quantification of the fault and event tree models using probability distributions for the various input quantities of interest (component failure rates, etc.) yields a probability distribution describing the uncertainty about the likelihood of an accident.

Consequence Analysis

For most types of engineered systems, consequence analysis (like exposure assessment for health risk assessment, as discussed above) typically is comprised of two distinct stages—one having to do with the migration of any hazardous materials released in the accident from sources to sinks, and another having to do with the consequences of those materials for public health and safety (using dose-response relationships as appropriate, as in health risk assessment). Relevant consequence measures might include: the structural response of a building to an impact or explosion; the costs associated with property damage, loss of use, and facility repair; the amount of hazardous material or energy released to the environment as a result of the accident; and the numbers of fatalities or other health effects.

Risk Characterization

Risk characterization focuses on how information is portrayed for decision-making. Results can be presented graphically in a variety of formats. For example, probability distributions are often used to display uncertainty about quantities such as the frequency of an accident (see Figure 2). For damage types that can take on differing levels of severity (e.g., repair cost, number of fatalities), a complementary cumulative distribution function can be used to show the frequency of exceeding any given damage level X (see Figure 3); uncertainty about such exceedance frequencies is sometimes represented by showing multiple complementary cumulative distributions in the same figure.

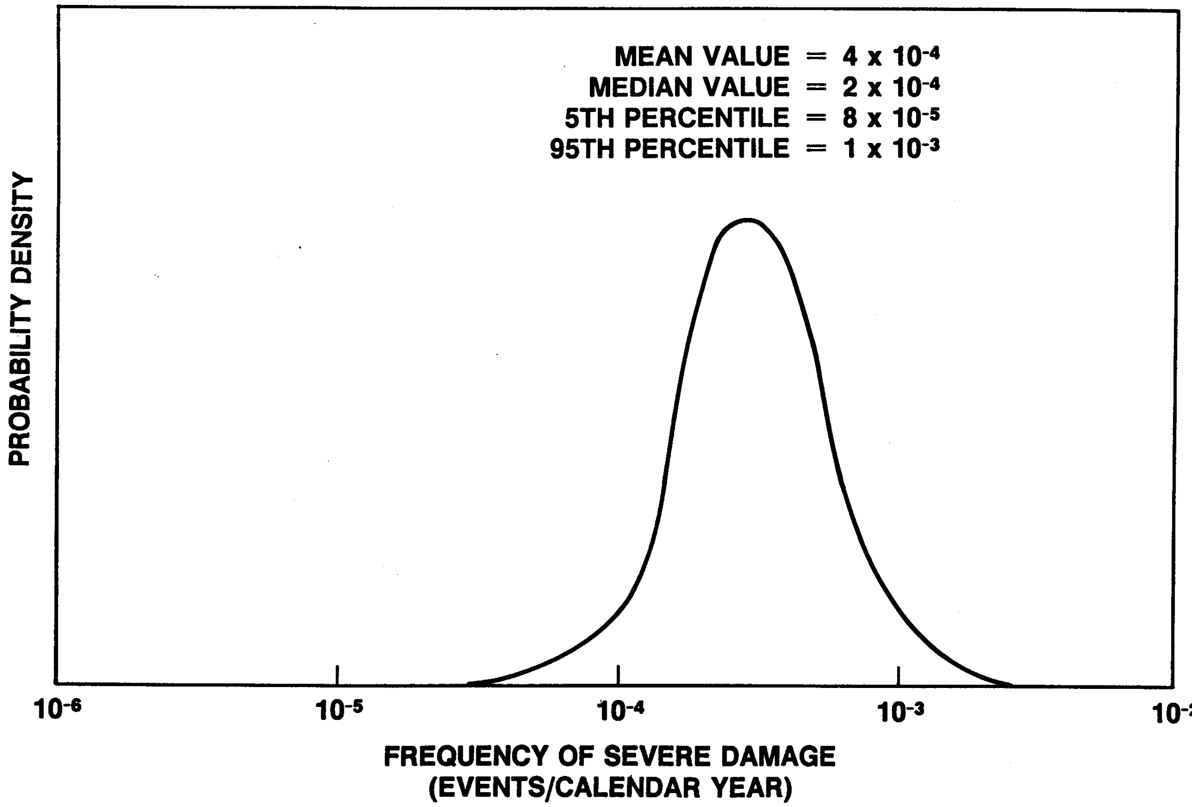

Figure 2. Probability Distribution for Frequency of Severe Damage

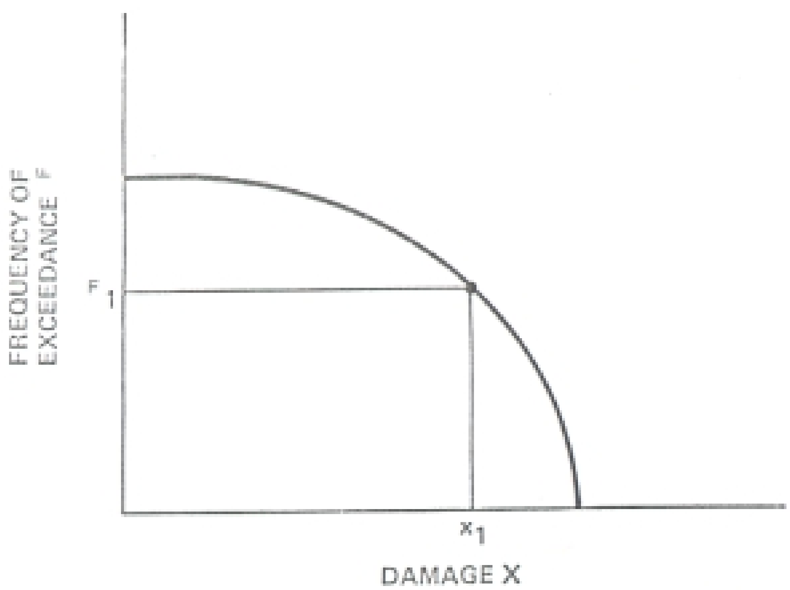

Figure 3. Complementary Cumulative Distribution

After the initial risk assessment has been completed, the model can then be used for risk management purposes—e.g., to identify the dominant contributors to risk, and assess the benefits of risk reduction options (such as design and procedure changes). The application of PRA has resulted in a number of examples of successful risk management; a few examples are given in Garrick (1987).

C. Spatial Dimensions as Organizing Principles for Health and Engineering Risk: The Relevance of Proximity

Proximity is a key factor primarily in the exposure portion of the risk equation. With respect to environmental and health applications, distance from a hazard can dilute or otherwise minimize risk, except if there are unusual amplifying or concentrating effects distant from the origins. Moreover, even estimated distances will depend upon how data on the location of human receptors and sources are defined. Despite the importance of spatial relationships, however, risk analyses rarely introduce sophisticated spatial concepts.

For example, engineering risk assessments do sometimes include spatial interactions analysis (Liming and Bennett 1994)—e.g., to assess how the failure of one component can affect the performance of other nearby components. However, many standard risk assessment references, such as Kumamoto and Henley (1996) and Modarres et al. (1999), do not even mention spatial interactions analysis explicitly. Even when risk assessments do include this type of analysis, the methodology for doing so tends to be relatively ad hoc, and has so far taken little or no advantage of the more sophisticated types of information that could be available using modern geographic information systems (GIS) and related technology.

Similarly, the proximity of a damaged or contaminated area to other important facilities could affect the perceived severity of particular scenarios. For example, if damage can occur at any of several locations, different facilities and exposed populations could be affected. The distribution of social and economic consequences can also vary considerably depending on the spatial distribution of impacts. This raises questions about geographical and environmental equity; for example, there has been a lot of dispute about how to specify spatial parameters, given the implication of the methods of specification for equity (Zimmerman 1994b).

Understanding and representing the spatial relationships among the various components of the system can dramatically alter the structure of the risk assessment model. It will also affect the conditional probabilities in the model, since failure probabilities in general are spatially dependent.

IV. Understanding Uncertainty

Uncertainty and risk are interconnected. Without uncertainty, risk loses much of its meaning. Estimates of risk resulting from risk assessments (especially for extreme or catastrophic events) will typically be highly uncertain, due in part to the sparse nature of the data on such events. Therefore, as indicated at the outset of this paper, uncertainty is central to the idea of risk. In fact, it is the uncertain and probabilistic nature of risk that requires risk assessment methodologies in the first place. As a result, uncertainty analysis (i.e., the mathematical methods for propagating uncertainties through risk models to quantify the uncertainty about the outcome quantity of interest; Bier 1987) is a significant part of risk analysis, and decision-making under uncertainty (the theory of how to take those uncertainties into account in an optimal manner in decisions) is an important part of risk management.

Morgan and Henrion (1990) present what is probably the best available taxonomy of types of uncertainty. In particular, they note that uncertainty can arise from:

- Statistical variation (e.g., "random error in direct measurements of a quantity");
- Systematic error (e.g., "biases in the measuring apparatus and experimental procedure");
- Subjective judgment (e.g., for quantities where empirical data is largely unavailable);
- Linguistic imprecision (e.g., translation of verbal phrases into numerical probabilities);
- Variability (e.g., quantities that vary over time or space, or from one person to another);
- Inherent randomness or unpredictability (which cannot be reduced by further research);
- Disagreement (e.g., among multiple experts);
- Approximation (e.g., due to limits in the spatial resolution of a model); and
- Uncertainty about the most appropriate model to represent some phenomenon.

The problem of decision-making under uncertainty is addressed by standard decision theory (see for example Raiffa 1968; Clemen 1991). Decision theory provides guidance not only on how to select the best of several uncertain options, but also on "the value of information"—i.e., whether it is preferable to make a final decision, or gather more information and defer the decision until that information becomes available. Thus, one important use of risk assessments (especially those that include an explicit statement of uncertainties) is to assist in determining whether additional research is needed.

State-of-knowledge uncertainty and population variability have different implications for decision-making. Thus, the National Research Council (1994) states: "uncertainty forces decision makers to judge how probable it is that risks will be overestimated or underestimated for every member of the exposed population, whereas variability forces them to cope with the certainty that different individuals will be subjected to risks both above and below any reference point one chooses." For example, if the various plants or facilities in a population differ a great deal from each other (i.e., high population variability), but the risk at each one is accurately known (i.e., low state of knowledge uncertainty), it makes sense to target risk reduction efforts at the facilities with the highest estimated risks. By contrast, when the plants or facilities in a population have relatively similar risk estimates (i.e., low variability), but those estimates are all quite uncertain, then the facility with the highest risk estimate may not actually have the highest risk. In this case, it may be better to spread one's risk reduction efforts across the entire population, to protect against the possibility that risks have been misestimated and a high-risk facility has been neglected. Similarly, in human health risk assessment with large population variability, it makes sense to target protection measures primarily aimed at the most vulnerable sub-populations. Large state of knowledge uncertainty suggests that further research may be desirable before making a decision on how or whether to implement risk reduction actions.

V. The Challenges of Predicting Human Perceptions, Behavior, and Performance in Extreme Events for Risk Assessment

The World Trade Center disaster and others that have been studied in the risk literature point out the central role of human performance in risk assessment. This shows up in two crucial areas—the importance of human behavior in responding to an event (e.g., in evacuations), and the challenges of assessing and managing intentional (as opposed to natural or accidental) hazards. Analyzing human behavior poses unique difficulties for risk analysis.

A. Predicting Human Response to Extreme Events

With regard to human response to severe events, evacuation responses in actual emergencies differ substantially from ideal performance in evacuation tests and simulations (see for example Turner and Pidgeon 1997; Fahy and Proux 1997). Similarly, it is known that the behavioral assumptions underlying many building codes and disaster-response strategies are flawed (Pauls 1998; Sime 1999). For instance, building codes regulate the distance to the nearest exit, even though people often use more distant but familiar exits (Sime 1999). Similarly, Turner and Pidgeon (1997) discuss the Summerland fire in 1973, noting: "the stereotyped view of the public and its likely behaviour in the case of a fire ignored those parents who were separated from their children…on another floor." Recent research results (e.g., Peponis et al., 1990; Willham 1992) provide the start of a more scientific understanding of how people move within buildings, but will need to be extended to provide a reliable basis for predicting human response.

Note also that human behavior is extremely variable. For example, Brannigan and colleagues (1998, 2001) have pointed out that a building designed to ensure safe evacuation of hypothetical "average healthy individuals" may be not adequate to allow evacuation of people who are elderly, ill, or disabled. Thus, even a building that had been "shown" to be safe based on a risk assessment could still become the site of a major disaster if a fire or other emergency took place at a time when individuals with characteristics different than those assumed in the assessment were in the building.

Familiarity with a particular environment can also vary widely, with possible adverse consequences in the event of an emergency. For example, Turner and Pidgeon (1997) cite the Hixon rail accident (which killed 11 people in 1968) and the Summerland fire (which killed 50 people in 1973), both in the United Kingdom, as examples in which either the likelihood or the consequences of the event were greatly exacerbated by the presence of "strangers" who were not familiar with the environment. In the case of the Hixon rail accident, truckers driving a slow-moving vehicle were not familiar with the appropriate procedures to use at a new rail crossing. In the case of the Summerland fire, the fact that the building was a recreational center used by visitors who were not familiar with the building and its evacuation routes posed a problem in the evacuation. Turner and Pidgeon (1997) conclude, "The Summerland fire highlights the special problems posed for those responsible for hazardous situations where safe operation relies to some extent upon the safe behaviour of strangers."

Historically, risk assessment had been viewed as much more applicable to errors of omission (e.g., neglecting to perform a step in a well-defined procedure) than to errors of commission (e.g., deliberately undertaking an action not specified in procedures); see for example Bier (1999). More recently, it has been recognized that even errors of commission are generally "rational actions" governed by an individual's perception of the situation at hand (Julius et al. 1995; Hsueh and Mosleh, 1996). This recognition makes errors of commission more amenable to assessment than was previously thought, and there have by now been several pilot studies incorporating errors of commission into risk assessments (see for example Julius et al. 1995). However, even those methods to date have focused primarily on predicting the behavior of well-trained operators in occupational settings—e.g., anesthesiologists (Pate-Cornell et al., 1996), nuclear power plant control room operators (Julius et al. 1995; Hsueh and Mosleh 1996). The emphasis on well-trained professionals substantially restricts the range of possible behavior that must be anticipated. Predicting the behavior and capabilities of members of the lay public, under a wide variety of possible emergency circumstances for which they may not be trained or prepared, will be a much more difficult challenge (Brannigan and Smidts 1998; Brannigan et al. 2001).

B. Intentional versus Unintentional Hazards

Predicting human behavior in emergency situations is already difficult. However, in attempting to estimate and manage the risks of intentional attacks, further difficulties become apparent. First, as pointed out by Woo (2002), "some idea of event likelihood is needed for intelligent benefit-cost analysis." However, estimating the likelihood and nature of intentional attacks is an area with which most risk assessors are not yet familiar, although there has been some related work on this problem in other fields. For example, Dickey (1980) interviewed bank robbers in custody to understand the criteria that they used in choosing banks to rob; he found that robbers preferred banks located near major highways, and banks with a single point in the lobby from which the robber could see all of the employees at once. Similarly, Crowe (1991) and de Becker (1997) report that criminals choose targets based not only on the attractiveness of the target, but also on the likelihood that they would be discovered and apprehended. Interviews with incarcerated terrorists could presumably be used to explore the criteria they use in selecting targets, to factor into quantitative risk assessments.

More significantly, protection from a knowledgeable and adaptable adversary is a fundamentally different challenge than protection against accidents or acts of nature. For example, earthquakes do not get stronger or smarter just because we have defended our buildings against them. However, if one's adversaries know or can easily learn about one's defensive measures, then they can actively choose to either bypass or circumvent those defenses. Progress in and increased reliance upon detection technologies has made this a more important possibility to take into account. For example, metal-screening devices had prior to September 11th increased the security and safety of air travel. A network news report early in 2002 suggested that the box cutters used by the terrorists on September 11th to gain control of the hijacked airplanes are said to have fallen just below the detection settings of such screening devices.

As noted by Dresher (1961), optimal allocation of defensive resources requires that "each of the defended targets yield the same payoff to the attacker." Thus, even if some components can be hardened quite inexpensively, focusing protective investments only on those can lead to wasted resources, if adversaries instead choose to attack targets that cannot be hardened as cost-effectively. In other words, critical assets must be defended against all possible attacks, which is much more difficult than just shoring up a few "weak links." As a result, Ravid (2001) has concluded that security improvements are generally more costly than safety improvements: "investment in defensive measures, unlike investment in safety measures, saves a lower number of lives (or other sort of damages) than the apparent direct contribution of those measures."

For the same reason, many defensive measures may be much less effective if they are widely known (or can be easily discovered by adversaries), since adversaries can then either design attacks to overcome those defenses, or else shift to less well-defended targets. However, due to the openness of our society, it may be quite difficult to keep some types of protective measures secret, since members of the public may demand to know how critical facilities are defended. This further increases the difficulty of defending high-profile targets that are of significant public concern. As just one example, the proposal to sterilize mail to protect against future anthrax attacks might be an effective defense if the installation of sterilization equipment could be kept secret. However, this would probably not be possible in our open society, especially given the outcry of public concern about the threat of anthrax spores being sent through the mail. If installation of anthrax sterilization equipment is public knowledge, it will likely only cause future attackers to find a different (and possibly more effective) means of delivery instead of the public mail. In that case, the proposed $40 million of sterilization equipment may never sterilize a single anthrax spore!

Similarly, the optimal strategy for reducing vulnerability to intentional hazards may depend critically on the assumptions made regarding attacker behavior and motivations. Consider the question of what attackers decide to do when a few "signature" buildings have been fortified to protect them against prospective attacks. If attackers simply move on to target progressively less well defended buildings, then expensive methods of protecting signature buildings may not actually be effective at protecting the public health and safety. Kunreuther and Heal (2002) point out a similar problem with installation of burglar alarms for theft protection, in which installation of an alarm at one house increases the risk of burglary at neighboring homes.

Thus, methods that are so expensive that they can be employed only in those few buildings that are perceived to be most "at risk" will protect people and assets in those particular buildings, but may yield little or no overall public health and safety benefit if attackers simply target other buildings. In that case, only methods that are inexpensive enough for widespread application, and/or mobile (such as emergency response capabilities), will substantially reduce overall levels of societal risk. By contrast, if attackers are interested only in the few highest-profile targets, then expensive methods of protecting those few targets may be potentially worthwhile from a societal as well as a private point of view. The fact that assumptions regarding attacker behavior and motivations are important in determining the expected effectiveness of protective measures illustrates the need for relevant social science research.

VI. Implications of the World Trade Center and Other Disasters for Learning from the Past

One challenge of the World Trade Center disaster is that it had been either unexpected or counterexpected. Thus, some risk management mechanisms that could have either reduced the likelihood of the event or minimized its consequences were not in place. As just one illustration, the insurance policy covering the World Trade Center had not anticipated that more than one building could be destroyed in a single event. As a result of the failure to anticipate this possibility and draft contract language to cover it, there was a significant dispute about whether the building owner was entitled to collect insurance on the destruction of the second building (SR International Business Insurance Co. Ltd. 2001). On the other hand, the configuration of many of the utility networks and their density, given the overall density of the area, provided the kind of resiliency, whether expected or not, to rebound relatively quickly (Zimmerman 2001).

Since September 11th, a number of past experiences have been put forth that could perhaps have helped to identify the risk of such an attack, or could help predict the likelihood of such events in the future. For example, Barnett (2001) states:

> After the disaster, it was widely (and truthfully) reported from the aviation community that no one there imagined an event like the one that happened. But could that circumstance reflect a failure of our imaginations?
>
> There were, after all, lots of events that could be interpreted as precursors of the calamity. In 1994, terrorists tried to hijack an Air France jet in Algiers, with the aim of using the plane to destroy the Eiffel Tower. (The plan was thwarted, but not without three passenger fatalities.) A disgruntled FedEx employee tried in 1998 to hijack one of the company's planes, intending to crash it into the company's Memphis headquarters. Hijackers had caused disasters on Ethiopian Airlines (1996) and in Mainland China (1994) that resulted in triple-digit death tolls (including themselves). Suicidal pilots apparently crashed two crowded jets in the late 1990s, with no survivors. And terrorists plotted in 1995 to destroy a dozen U.S. jets coming home from Asia. Their plans were foiled literally at the last minute: Having successfully exploded a small test bomb on a Philippine Airlines plane, some of them were in line to board an U.S.-bound jet in Bangkok when the conspiracy was uncovered.
>
> Thus, all the elements of the Sept. 11 catastrophe—the idea of using planes as weapons, suicidal individuals in the cockpit, and a willingness to take thousands of innocent lives—had historical precedent.

In addition to the examples cited by Barnett, a B-25 bomber crashed into the Empire State Building in 1945, resulting in the death of 14 people and many more injured (Mandell 2001), and a helicopter hit the Pan Am building (now the Met Life building) in 1977, resulting in five fatalities (Horsley 2002). However, both events were accidents, not acts of terrorism. The planes were also smaller, traveling at much lower speeds than what have now been estimated as the speeds of the airplanes that attacked the World Trade Center towers, and probably not loaded with as much fuel. Perhaps as a result, the structures were more resilient to the levels of impact that they experienced. Yet another past analogy that has been suggested is the attack on Pearl Harbor, the only previous attack on U.S. soil within the last hundred years.

Perhaps the most direct analogy that has been cited is the bombing of the World Trade Center on February 26, 1993. The 1993 bombing of the World Trade Center led to a number of risk reduction measures dealing with various aspects of disaster management. These included the establishment of the Urban Search and Rescue teams in the mid-1990s, and the establishment of the Center for National Response that opened in December 2000, which includes training in a collapsed subway tunnel (Gittrich 2002) to increase the likelihood that responders will be able to find victims, and thus to lower the risk of death immediately following catastrophic events.

For natural hazards, similar examples exist where the occurrence of a catastrophic event could perhaps have been predicted from events immediately preceding it. For example, the severity of the Midwest floods in 1993 could have been anticipated from the gradual buildup of soil moisture to the point of saturation in the months preceding the floods (see references cited in Zimmerman 1994a).

While it would of course be a mistake to demand perfect foresight on the part of risk managers, the fact that past experience foreshadowed the possibility of the World Trade Center and other disasters points out the need for good methods of learning from past experience. March et al. (1991) discuss numerous strategies for “learning from samples of one or fewer,” which is often necessary in order to successfully avoid severe events. The general strategies they suggest include ways of viewing past history more “richly” (to extract as much information as possible from unique events), and using near misses (such as those described above) and hypothesized scenarios to enrich the available database.

Collection, reporting, and analysis of data on near misses or “precursors” to severe events (Bier 1998) are important steps to facilitate learning from past experience. For example, van der Schaaf (1992) notes that reporting of near misses is valuable for several purposes: (1) “to gain qualitative insight into how (small) failures or errors develop into near misses” or more serious events; (2) “to arrive at a statistically reliable quantitative insight into the occurrence of factors or combinations of factors giving rise to incidents”; and (3) “to maintain a certain level of alertness to danger” and avoid complacency. For that reason, William Corcoran, a root cause analysis consultant, recommends that root cause analyses of undesired events should always include an analysis of how the observed event could have been made worse than it was.

The analysis of NASA’s response to precursors to the Challenger disaster by Vaughan (1996) also shows the importance of being vigilant to incipient signs of problems, rather than ignoring or rationalizing them and minimizing the attendant dangers. Thus, for example, Turner and Pidgeon (1997) note that some past disasters occurred in part because warnings of trouble were dismissed as coming from “outsiders”—e.g., members of the lay public. Similarly, sociologist Ron Westrum notes that “generative” organizations actively seek out information, regardless of its source, rather than dismissing or compartmentalizing it.

In connecting information from prior events more closely to the process of risk assessment, it is helpful to look at the occurrence frequency and consequence analysis components of risk assessment separately. As described earlier, risk is a combination of the probability of occurrence of an event and the severity of the consequences. Uncertainty reduction depends on the extent to which we can predict both occurrences and consequences, and the use of past experiences is one basis for such predictions.

Methods of precursor analysis (Bier 1998) provide one strategy for using prior event information in estimating the frequencies of extreme events. Other methods have been discussed in Bier et al. (1999).

While perhaps less well recognized, analogies from prior events are also needed to estimate health and safety consequences, just as much as to estimate occurrence frequencies, since modeling or predicting the consequences of an event requires that we have some basis for predicting the particular conditions that are likely to occur afterwards. Thus, as New York City moved away from the initial crisis of survival at the World Trade Center site, there have been burgeoning concerns over air quality and health, and the search for precedents or analogies to help address this issue has been intense. In particular, concentrated exposure to particulate matter was a key consequence of the collapse of the buildings. Here some of the intense exposures to particulates that occurred as a result of natural disasters are instructive. The eruption of Mount St. Helens, for example, might be an important model for understanding the nature of the particulate releases and exposures following the collapse of the World Trade Center buildings. The understanding of how particulate matter behaves and its effects on human health also benefited considerably from at least two major documents that had been completed just prior to September 11th: the Criteria Document for Particulate Matter prepared under the Clean Air Act (U.S. EPA, ORD, October 1999); and a series of three volumes produced by the National Research Council entitled *Research Priorities for Airborne Particulate Matter* (National Research Council 1998, 1999, and 2001). However, regardless of the ability to predict environmental and health consequences from previous events and the gains from previous scientific knowledge, the extent to which exposure assessments can reflect the range of individual circumstances and the extent to which preexisting standards are adequate in light of previously unanticipated exposures have become serious questions.

VII. Conclusions Regarding the Relationship of Risk Assessment to Extreme Events

Risk assessment in its current form (as a systems-oriented method that is flexible enough to handle a variety of alternative conditions) is a vital tool for dealing with extreme events, for both health and the built environment. However, the capabilities of traditional risk assessment are challenged when we attempt to apply it to extreme and unanticipated events of the nature and magnitude of the September 11th disaster (Bier et al. 1999).

For example, events of extreme, unexpected, and catastrophic proportions, such as the destruction of the World Trade Center, point out the need for methodological improvements in risk assessment to improve and more fully incorporate spatial dimensions. Spatial dimensions directly relate the risk to the exposed populations, and are critical since such relationships can dramatically alter the consequences of an event. Another aspect of the spatial dimension is the interconnections and interdependencies among support systems that must be taken into account in risk assessment, since these relationships can also dramatically affect event consequences.

Another challenge to risk assessment is the fact that many human considerations—values, attitudes, beliefs, and behavior—are critical inputs throughout a risk assessment and must be part of its formulation, not just add-ons.

Acknowledgments

This material is based upon work supported in part by the U.S. Army Research Laboratory and the U.S. Army Research Office under grant number DAAD19-01-1-0502. This work is also supported in part by the Institute for Civil Infrastructure Systems (ICIS) at New York University (in partnership with Cornell University, Polytechnic University of New York, and the University of Southern California) funded by the National Science Foundation under Cooperative Agreement No. CMS-9728805. Any opinions, findings, and conclusions or recommendations expressed in this document are those of the author(s) and do not necessarily reflect the views of the National Science Foundation or other sponsors.

References

A. Barnett, "The Worst Day Ever: The Sept. 11 Catastrophe Raises Many Troubling Questions," *OR/MS Today*, December 2001, http://lionhrtpub.com/orms/orms-12-01/worstdayfr.html.

V. M. Bier, "Uncertainty Analysis as Applied to Probabilistic Risk Assessment," in V. T. Covello, L. B. Lave, A. Moghissi, and V. R. R. Uppuluri (editors), *Uncertainty in Risk Assessment, Risk Management, and Decision Making.* New York, NY: Plenum, 1987, pp. 469-478

V. M. Bier, "An Overview of Probabilistic Risk Analysis for Complex Engineered Systems," in V. Molak (editor), *Fundamentals of Risk Analysis and Risk Management.* Boca Raton, FL: CRC Press, 1997.

V. M. Bier, *Accident Sequence Precursors and Probabilistic Risk Analysis.* University of Maryland, College Park, Maryland, 1998.

V. M. Bier, "Challenges to the Acceptance of Probabilistic Risk Analysis," *Risk Analysis*, Volume 19, 1999, pp. 703-710.

V. M. Bier, "On the State of the Art: Risk Communication to Decision Makers," *Reliability Engineering and System Safety*, Volume 71, 2001, pp. 151-157.

V. M. Bier, Y. Y. Haimes, J. H. Lambert, N. C. Matalas, and R. Zimmerman, "A Survey of Approaches for Assessing and Managing the Risk of Extremes," *Risk Analysis*, Volume 19, 1999, pp. 83-94.

D. Blockley, *Engineering Safety*. London, UK: McGraw-Hill, 1992.

V. Brannigan and C. Smidts, "Performance Based Fire Safety Regulation under Intentional Uncertainty," *Proceedings of the First International Symposium on Human Behavior and Fire*, University of Ulster, 1998, pages 411-420. Reprinted in *Fire and Materials*.

V. Brannigan, C. Smidts, and A. Kilpatrick, "Human Behavior and Risk Based Fire Regulation, 2nd Inernational Symposium on Human Behavior and Fire, Massachusetts Institute of Technology, 2001, pp. 411-418.

R. T. Clemen and T. Reilly, *Making Hard Decisions with Decision Tools*. Pacific Grove, CA: Duxbury, 2001.

T. D. Crowe, *Crime Prevention through Environmental Design: Applications of Architectural Design and Space Management Concepts.* Boston, MA: Butterworth-Heinemann, 2000.

A. C. Cullen and H. C. Frey, *Probabilistic Techniques in Exposure Assessment: A Handbook for Dealing with Variability and Uncertainty in Models and Inputs.* New York, NY: Plenum, 1999.

G. de Becker, *The Gift of Fear: Survival Signals that Protect us from Violence.* Boston, MA: Little, Brown and Company, 1997.

T. S. Dickey, *Bank Robbery—Architectural Implications from the Criminal's Point of View*, unpublished Master's Thesis, Georgia Institute of Technology College of Architecture, 1980.

M. Dresher, *Games of Strategy: Theory and Applications.* Englewood Cliffs, NJ: Prentice-Hall, 1961.

R. F. Fahy and G. Proulx, "Human Behavior in the World Trade Center Evacuation," *Proceedings of the 5th International Symposium of the International Association for Fire Safety Science*, Boston, 1997, pp. 713-724.

B. J. Garrick, "Examining the Realities of Risk Management," in V. T. Covello, L. B. Lave, A. Moghissi, and V. R. R. Uppuluri (editors), *Uncertainty in Risk Assessment, Risk Management, and Decision Making.* New York, NY: Plenum, 1987.

G. Gittrich, "Where They Learn to Fight Terror," *Daily News*, February 17, 2002.

Y. Y. Haimes, *Risk Modeling, Assessment, and Management.* New York, NY: Wiley, 1998.

Y. Y. Haimes and C. Schneiter, "Covey's Seven Habits and the Systems Approach," *IEEE Transactions on Systems, Man, and Cybernetics*, Volume 26, 1996, pp. 483-487.

C. B. Horsley, "The Metlife Building," http://www.thecityreview.com/panam.html (accessed March 19, 2002).

K.-S. Hsueh and A. M. Mosleh, "The Development and Application of the Accident Dynamic Simulator for Dynamic Probabilistic Risk Assessment of Nuclear Power Plants," *Reliability Engineering and System Safety*, Volume 52, 1996, pp. 297-314.

J. Julius, E. Jorgenson, G. W. Parry, and A. M. Mosleh, "A Procedure for the Analysis of Errors of Commission in a Probabilistic Safety Assessment of a Nuclear Power Plant at Full Power," *Reliability Engineering and System Safety*, Volume 50, 1995, pp. 189-201.

S. Kaplan, "On a 'Two-Stage' Bayesian Procedure for Determining Failure Rates from Experiential Data," *IEEE Transactions on Power Apparatus and Systems*, Volume PAS-102, 1983, pp. 195-202.

S. Kaplan and B. J. Garrick, "On the Quantitative Definition of Risk," *Risk Analysis*, Volume 1, 1981, pp. 11-27.

H. Kumamoto and E. J. Henley, *Probabilistic Risk Assessment and Management for Engineers and Scientists*. New York, NY: IEEE Press, 1996.

H. Kunreuther and G. Heal, "Interdependent Security: The Case of Identical Agents," paper to be presented at the National Bureau of Economic Research, Inc., Insurance Project Workshop, February 1, 2002, http://www.nber.org/~confer/2002/insw02/kunreuther.pdf.

L. B. Lave and T. Balvanyos, "Risk Analysis and Management of Dam Safety," *Risk Analysis*, Volume 18, 1998, pp. 455-462.

J. K. Liming and L. A. Bennett, "Advances in the Methodology for the Analysis of Location-Dependent Hazards for Probabilistic Risk Assessment (PRA)," in *Proceedings of PSAM-II: International Conference on Probabilistic Safety Assessment and Management*, March 20-25, San Diego, California, 1994.

J. Mandell, "Looking To History," *Gotham Gazette*, October 1, 2001, http://www.gothamgazette.com/commentary/103.mandell.shtml.

J. G. March, L. S. Sproull, and M. Tamuz, "Learning from Samples of One or Fewer," *Organization Science*, Volume 2, 1991, pp. 1-13.

H. F. Martz and R. A. Waller, *Bayesian Reliability Analysis.* New York, NY: Wiley, 1982.

N. J. McCormick, *Reliability and Risk Analysis: Methods and Nuclear Power Applications*. San Diego, CA: Academic Press, 1981.

Merriam-Webster's Collegiate Dictionary, 2002, http://www.m-w.com/home.htm.

Dennis S. Mileti, *Disasters by Design*. Washington, DC: Joseph Henry Press, 1999.

M. Modarres, M. Kaminskiy, and V. Krivtsov, *Reliability Engineering and Risk Analysis*. New York, NY: Marcel Dekker, 1999.

M. G. Morgan and M. Henrion, *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis.* Cambridge, MA: Cambridge University Press, 1990.

F. Moses, "Probabilistic-Based Structural Specifications," *Risk Analysis*, Volume 18, 1998, pp. 445-454.

National Research Council, Committee on Institutional Means for Assessment of Risks to Public Health, *Risk Assessment in the Federal Government: Managing the Process.* Washington, DC: National Academy Press, 1983.

National Research Council, Committee on Risk Assessment of Hazardous Air Pollutants, *Science and Judgment in Risk Assessment.* Washington, DC: National Academy Press, 1994.

National Research Council, *Research Priorities for Airborne Particulate Matter I: Immediate Priorities and a Long-Range Research Portfolio*. Washington, DC: National Academy Press, 1998.

National Research Council, *Research Priorities for Airborne Particulate Matter II: Evaluating Research Progress and Updating the Portfolio*. Washington, DC: National Academy Press, 1999.

National Research Council, *Research Priorities for Airborne Particulate Matter III: Early Research Progress*. Washington, DC: National Academy Press, 2001.

National Research Council, *Cybersecurity Today and Tomorrow: Pay Now or Pay Later*. Washington, DC: National Academy Press, 2002.

E. Pate-Cornell, "Fault Trees vs. Event Trees in Reliability Analysis," *Risk Analysis*, Volume 4, pp. 177-186, 1984.

E. Pate-Cornell, D. Murphy, L. Lakats, and D. Gaba, "Patient Risk in Anesthesia: Probabilistic Risk Analysis and Management Improvements," *Annals of Operations Research*, Volume 67, 1996, pp. 211-233.

J. Pauls, "A Personal Perspective on Research, Consulting and Codes/Standards in Fire-Related Human Behavior, 1969-1997, with an Emphasis on Space and Time Factors," presented at the First International Symposium on Human Behavior in Fires, Jordanstown, Northern Ireland, 1998.

J. Peponis, C. Zimring, and Y. K. Choi, "Finding the Building in Wayfinding," *Environment and Behavior*, Volume 22, 1990, pp. 555-590.

The Presidential/Congressional Commission on Risk Assessment and Risk Management, *Risk Assessment and Risk Management in Regulatory Decision-Making*. Final Report, two volumes. Washington, DC: The Commission, 1997.

H. Raiffa, *Decision Analysis: Introductory Lectures on Choices under Uncertainty*. Reading, Massachusetts: Addison-Wesley, 1968.

I. Ravid, “Theater Ballistic Missiles and Asymmetric War,” draft, 2001.

W. D. Rowe, *An Anatomy of Risk*. New York, NY: Wiley, 1977.

A. P. Sage and E. B. White, **“**Methodologies for Risk and Hazard Asessment: A Survey and Status Report,**”** *IEEE Transactions on Systems, Man and Cybernetics*, SMC-10(8), pp. 425-445.

G. L. S. Shackle, *Imagination and the Nature of Choice.* Edinburgh, UK: Edinburgh University Press, 1972.

J. Sime, “Understanding Human Behavior in Fires: An Emerging Theory of Occupancy,” presented at the University of Ulster Faculty of Engineering, School of the Built Environment, Fire Safety Engineering Research and Technology Marks and Spencer Guest Lecture, Ulster, 1999.

SR International Business Insurance Co. Ltd., Plaintiff, v. World Trade Center Properties LLC et al., Defendants, “Complaint for Declaratory Relief,” filed in the United States District Court for the Southern District of New York, October 27, 2001.

P. C. Stern and H. V. Fineberg, editors, *Understanding Risk: Informing Decisions in a Democratic Society*. Washington, DC: National Academy Press, 1996.

B. A. Turner and N. F. Pidgeon, *Man-Made Disasters*, second edition. Oxford, UK: Butterworth-Heinemann, 1997.

U.S. Environmental Protection Agency (EPA), *Exposure Factors Handbook* (Final), Volumes I, II, III, 1999.

U.S. EPA, *Risk Assessment Guidance for Superfund Volume I. Human Health Evaluation Manual (Part A)*. Interim Final Report. Washington, DC: U.S. EPA, December 1989.

U.S. EPA, Office of Research and Development (ORD), *Air Quality Criteria for Particulate Matter*. Washington, DC: U.S. EPA, October 1999.

U.S. EPA, ORD, *Strategic Plan for the Office of Research and Development,* EPA/600/R-96/059. Washington, DC: U.S. EPA, May 1996.

U.S. Nuclear Regulatory Commission, *Reactor Safety Study*, WASH-1400. Washington, DC: Nuclear Regulatory Commission, 1975.

T. W. van der Schaaf, *Near Miss Reporting in the Chemical Process Industry*. Eindhoven, Netherlands: Technical University of Eindhoven, 1992.

D. Vaughan, *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA.* Chicago, IL: University of Chicago Press, 1996.

D. B. Willham, *The Topological Properties of Wayfinding in Architecture*. 1992.
G. Woo, "Benefit-Cost Analyses for Malevolent Human Actions," note prepared for the Columbia/Penn Roundtable, http://www.ldeo.columbia.edu/CHRR/Roundtable/gwpalis.pdf.

R. Zimmerman, "After the Deluge," *The Sciences*, Volume 34 No. 4 (July/August 1994a), pp. 18-23.

R. Zimmerman, "Issues of Classification in Environmental Equity: How We Manage is How We Measure," *Fordham Urban Law Journal*, Volume 21, Spring 1994b, pp. 633-669.

R. Zimmerman, "Historical and Future Perspectives on Risk Perception and Communication," in *Risk Research and Management in Asian Perspective: Proceedings of the First China-Japan Conference on Risk Assessment and Management*, November 23-26, 1998, Beijing, China. Edited by Beijing Normal University, Society for Risk Analysis-Japan Section, Department of Earth Sciences-National Natural Science Foundation of China. International Academic Publishers, 1998, pp. 481-487.

R. Zimmerman, "Social Implications of Infrastructure Network Interactions," *Journal of Urban Technology*, Volume 8, No. 3 (December 2001), pp. 97-119.