1

Introduction

I began by trying to quantify technical risks, thinking
that if they were "put into perspective" through
comparison with familiar risks we could better judge
their social acceptability. I am ashamed
now of my naiveté, although I have the excuse
that this was more than twenty years ago,
while some people are still doing it today.

Harry Otway, 1992

Defining Risk

What is risk? What are the tools and methods used to evaluate particular health, environmental, technological, and other risks, and what are the limitations, uncertainties, and biases in these methods? How can and will the results found using those methods be used by individuals and groups?

This book is about modeling and calculating a variety of risks, understanding what we're trying to calculate, and why we would want to do so. First, however, what is risk? A simple, albeit "technocratic," definition of that risk is the probability that an outcome will occur times the consequence, or level of impact, should that outcome occur. To many people, risk suggests adverse outcomes; however, technical approaches to evaluating probabilities and outcomes are not limited to negative impacts. Rather, they represent positive or negative changes in state.

We can quantify risks in a number of ways, and often with considerable precision. While this quantification can be a useful tool, it is not the whole story. This book leads through technical and analytic methods used to evaluate and test risk, and then into the more intricate world of social valuation and decision theory to which Otway alludes. We begin with an exploration of the quantitative methods, and then expand the sphere of analysis to include uncertainty, economic, political, and social dimensions of risk understanding and management. Our operating principle is that when we can better understand and describe values (that is, what the outcomes and probabilities are likely to be and how complete our understanding is), we can make better decisions.

Sheila Jasanoff proposes that the role of risk assessment is to "offer a principled way of organizing what we know about the world, particularly about its weak spots and creaky joints" (Jasanoff 1993). In keeping with this philosophy, the goal of this book is not to produce "technocrats" who will apply these tools to decisions outside of a social context. Rather, we hope that our readers will learn not only how to "crank the numbers," but when and why they should, and how the numbers will be interpreted in a broader cultural context. Ideally, risk analysis responds to the needs of interested and affected groups and individuals; it is intended to inform, but not determine, decisions.

Examples of the pressing need for better risk analysis abound. At the microdecision level, this agenda includes evaluating the impacts of and possible responses to rare but potentially "catastrophic" risks; identifying mechanisms of disease (and consequently improving opportunities to cure or avoid them); comparing similar remedies to a single adverse situation; and evaluating the possibly different responses of adults and children to a potential risk factor.

This book introduces a diverse audience to the fundamental theories and methods for modeling and analyzing risk. As a synthetic approach to both the subject of risk and the standard risk analysis "tool kit" we envision the potential for wide use in the fields of environmental science, engineering risk/fault analysis, public policy and management, and science policy. In particular, these methods should be of interest to policy makers at the local, state, and federal level who are now confronted with legislation that requires them to perform risk and cost/benefit analyses prior to a range of actions.

Increasingly, professional decision makers such as engineers, environmental scientists, "policy wonks," and others find that they need to answer risk questions. They may be asked to generate a report on risks, or to recreate and critique how someone else created a report. They may need to be able to communicate their work to a skeptical public, or to a busy politician. They are also likely to find that they lack the tools to deal with these issues as they arise.

At the same time, the uninitiated are likely to see the process of risk assessment as enormously complex and problem specific. Looking at a single problem too closely can lead to two unsatisfactory end points.

One is to leave the problems "up to the experts," taking the results from the risk assessment "black box" at face value. The other is to get lost in the details of the problem at hand. This is unfortunate, since a few general tools can equip analysts to tackle most, if not all, problems of risk.

The fields of science and technology policy and environmental studies have only a limited number of unifying methods. The goal of our work is to develop a practical approach to formulating, solving, and then generalizing the theory and methods of risk analysis. This book provides a set of tools to clarify and define these methods, producing more than the current set of fascinating, but idiosyncratic and anecdotal, case studies. We seek to bridge the gap between qualitative "discussion" books, which provide little analytic or practical training; advanced modeling books and journal papers, which generally assume considerable prior knowledge on the part of the reader; and highly specialized works in the areas of medical epidemiology or industrial emissions. To do so, we present and suggest solutions to real-world problems using a variety of risk analysis methods.

The case studies we present include subjects as diverse as the health impacts of radon, trends in commercial and military flight safety, extrapolation from high-dose laboratory animal studies to low-dose human exposures, and some key decisions relevant to the proposed national high-level nuclear waste storage facility at Yucca Mountain, Nevada. The solutions to the exercises provide a springboard to the broader applications of each method to other technological, environmental, public health, and safety risk issues, as well as to forecasting and uncertainty. Additional unsolved problems reinforce the presentation. The methods include the scientific and quantitative methods used to evaluate risks, as well as analytical tools for social/political management and decision making.

The central theory and methods of risk covered in this book include order-of-magnitude estimation; cause-effect (especially dose-response) calculations; exposure assessment; extrapolations between experimental data and conditions relevant to the case being addressed; modeling and its limitations; fault-tree analysis; and managing and estimating uncertainty. While not the central focus of this book, statistics play a key role as a basic tool. We cover basic and intermediate statistics in chapter 3. Probabilistic risk assessment (PRA) methods, Bayesian analysis, and various techniques of uncertainty and forecast evaluation are presented and used throughout the book.

Note that we do not address the expanding field of financial risk. While many of the models and techniques are similar to those presented here, there is an entire literature devoted to that subject.

Structure of the Book

The goal of this book is to introduce the student to advanced risk analysis tools, but we believe that the risk analyst must be able to walk before she can run. In other words, gaining proficiency in the fundamentals of risk analysis necessarily precedes deeper understanding, and even mastering the basics can substantially aid decision making. Consequently, most of the book is directed at learning to manipulate various individual tools, and understanding their applications and limitations. Toward the end of the book we provide examples of real-world applications ranging from local, specific, and clearly definable risks to some that involve multiple stakeholders and substantial uncertainty. The remainder of this introductory chapter discusses the history of the risk policy process, the current status of risk analysis as a central but often ad hoc technique, and the main areas of agreement and dispute about definitions and methods.

The first section (chapters 2–4) covers the basic “tools of the trade.” Chapter 2 presents basic modeling techniques, both with and without numbers. The use of “stock and flow” models as an approach to identifying and quantifying exposures is presented first, followed by a number of models and techniques for quantifying cause-effect relationships.

Chapter 3 reviews the basic statistical techniques most commonly used in risk assessment. In general, solving the problems in this book requires fluency in high-school mathematics and basic statistics. For some problems calculus is a useful, although not necessary, prerequisite. (In fact, given the extent of uncertainty involved in many risk decisions, it should become clear that *over*-analysis can be a real problem.) While some of the models are easier to manipulate using more advanced mathematics, all the concepts and much of the implementation should be within the grasp of most college students. Many of the problems in this book have been used in the Princeton University graduate course

"Methods in Science and Technology Policy" (WWS-589), and have been taught without reference to calculus.

The beginning of the statistics chapter, designed more as a text than the rest of the book, is intended to be a review for those whose statistics are rusty; for the novice, a basic statistics class or text is recommended. The fourth chapter concludes the basic tools section with a discussion of variability, uncertainty, and forecasting, and provides two sophisticated statistical tools for dealing with variability and uncertainty: Bayesian analysis and probabilistic (Monte Carlo) analysis.

The second section of the book (chapters 5–8) applies these techniques to four important risk methodologies: structural models (e.g., toxicology), empirical models (e.g., epidemiology), exposure assessment, and technological risk assessment. Many of the problems address environmental risk, simply because that is where the authors have the most experience. However, a range of other issues are included, as well as discussion of how these methods can be applied in other fields.

The final section (chapters 9 and 10) deals with social aspects of risk: how people perceive risks, how people learn and communicate about risk, and how risk assessment can be incorporated into private and public decisions. The ninth chapter reiterates that the *application* of these tools should be limited to, motivated by, and designed to inform stakeholder and policy needs. This chapter puts the rest of the book into the decision-making context, introducing and critiquing some formal methods for both comparing among diverse risks and incorporating diverse interests. The final chapter discusses the human agent, and how perceptions of risk by both experts and nonexperts, as well as risk communication methods, influence risk decisions.

Risk analysis and computers complement one another very well, and most risk classes we are aware of incorporate a variety of software packages. Several of the problems in this book require the use of spreadsheets and risk software. In writing the problems, the authors generally used Microsoft Excel and the Crystal Ball and solver.xls add-ins, but other packages (such as Stella and @Risk) are of course acceptable.

Even small risk decisions may require many steps. No single problem or chapter can make the reader a "fully qualified risk analyst," but as a whole this book should enable the reader to synthesize the individual steps, combining them into coherent decisions. It will also promote enough healthy skepticism to guard against blind faith in any single methodology.

A book like this can never be truly "final." New solutions to old problems may be proposed by the readers, new information may change an existing problem, and emerging risks suggest novel methods and exercises. To keep pace, we are maintaining a website for this book at http://socrates.berkeley.edu/erg/swri. At the site, you will find

- Updated versions of problems and solutions in the book, including downloadable data files
- Copies of supplemental problems
- Solutions to supplemental problems, available to registered course instructors (if you are teaching from this book, contact the authors for a password)
- A dialog box to comment, append, or correct existing problems
- A dialog box to enter new problems and/or solutions

Our hope is that readers will contribute new cases that we will make available both on the World Wide Web site and in future editions of the book (with full attribution) as the fields of risk analysis and policy evolve.

Risk Analysis and Public Policy

In the past several decades, formal risk analysis has played an increasingly influential role in public policy, from the community to the international level. Although its outputs and uses are often (even usually) contentious, it has become a dominant tool for energy, environmental, health, and safety decisions, both public and private. More recently, risk analysis and cost-benefit analysis have been suggested by some (and even debated in Congress) as the *principal* tools for major federal environmental decisions, while others argue that the two methods have been oversold. While critiques abound, few scholars and practitioners would dispute the notion that an understanding of some essential tools of the trade is invaluable.

Risk analysis in one form or another has been used for centuries (see box 1-1, taken from Covello and Mumpower 1985). In the early 1970s, as risk analysis evolved into a major policy decision tool, Alvin Weinberg (1972) proposed that it falls into a special category of "trans-science . . . questions which can be asked of science, yet which cannot be

Box 1-1. Some historical highlights in risk analysis

About 3200 B.C.: The Asipu, a group of priests in the Tigris–Euphrates Valley establish a methodology:

- Hazard identification
- Generation of alternatives
- Data collection* and analysis
- Report creation

*Note that "data" included signs from the gods!

Arnobius, 4th century A.D., came up with decision analysis and first used the *dominance principle*, whereby a single option may be clearly superior to all others considered. Arnobius concluded that believing in God is a better choice than not believing, whether or not God actually exists. Note that Arnobius did not consider the possibility that a different God exists.

		State of nature	
		God exists	No God
Alternative	Believe	Good outcome (heaven)	Neutral outcome
	Don't believe	Bad outcome (hell)	Neutral outcome

King Edward II had to deal with the problem of smoke in London:

1285: Established a commission to study the problem.
1298: Commission called for voluntary reductions in use of soft coal.
1307: Royal proclamation banned soft coal, followed by a second commission to study why the proclamation was not being followed.

answered by science." Individuals and society need to make decisions on issues for which there are no certain outcomes, only probabilities, often highly uncertain.

Due to the "trans-scientific" nature of risk analysis, there will always be disputes about methods, end points, and models. Individual and societal values may not be separable from the quantitative analysis, determining what we choose to analyze. Tension over the use of quantitative analysis will be amplified by distributions of gains and losses, as well as prior commitments. Key goals of the risk analyst include extracting the good data from the bad, deciding which model best fits both the data and the underlying process, as well as understanding the limitations of available methods.

In some ways, risk analysis is a mature field, and a number of methods and techniques have become institutionalized. Yet in many profound ways, risk analysis remains immature. To some, the subject amounts to many fascinating case studies in search of a paradigm! The risks of contracting human immunodeficiency virus, of acquiring cancer from pesticides, of nuclear accidents, or of space shuttle disasters are regarded as important but idiosyncratic cases. To the extent that generalized lessons are not learned, science, technology, and environmental policy research has yet to find a common language of expression and analysis.

Despite a number of attempts to rationalize the use of risk analysis in the policy process, its role continues to be controversial. A 1983 National Resarch Council[1] (NRC) project, *Risk Assessment in the Federal Government*: *Managing the Process*, generally referred to as the "Red Book," sought to establish a risk assessment paradigm in the environmental context. It envisioned a sequence of Hazard Identification, followed by parallel Exposure and Dose-Response Evaluations, which are then combined to generate a Risk Characterization. Under this paradigm, once the hazard has been characterized, it can be used to inform risk management.

This approach embodies a technocratic philosophy promoting quantitative risk analysis as *the* solution to arbitrary and "irrational" risk policy decisions. Before he became a Supreme Court Justice, Stephen Breyer wrote in *Breaking the Vicious Circle* that risk assessors should be given an insulated, semi-autonomous decision-making role. John Graham, director of the Harvard Center for Risk Analysis, has campaigned similarly for rigorous training of risk assessors and a central federal department for risk assessment. Legislation that would have mandated quantitative risk assessment for all federal environmental, health, and safety regulations came close to being passed three times in the 1990s:

[1]The National Research Council is the research wing of the National Academy of Sciences.

SB110 in 1992, HR9 in 1994, and the Johnston–Robb Bill in 1996. One notable law that did pass eliminated the long-standing Delaney Clause, which had prohibited any known carcinogen as a food additive, regardless of the magnitude of the risk posed by that carcinogen. Under the new law, some levels of carcinogens may be acceptable.

Proponents of a participatory philosophy argue that risk analysis remains too subjective, and its implications too dependent on social context, to permit its removal from the public arena. Since decisions about values and preferences are made not just at the final decision stages but throughout the process, risk analysis necessarily combines both technical expertise *and* value choices. The implications of this interplay range from the inadvertent, as analysts make choices they believe are best without input from interested parties, to the antidemocratic, when the value decisions as well as the number crunching are intentionally restricted to a select group with a particular agenda.

While the "Red Book" approach has come to dominate the way the U.S. Environmental Protection Agency (EPA) approaches risk assessment, many feel that the firewall between "assessment" and "management" is artificial and distortional. Subsequent studies by the NRC (1994 and 1996) begin to address this issue. *Understanding Risk* (NRC 1996) identifies three "outstanding issues": inadequate analytical techniques, fundamental and continuing uncertainty, and a basic misconception of risk characterization. The study concludes that risk analysis must be decision driven and part of a process based upon mutual and recursive analytic-deliberative efforts involving all "interested and affected parties." While clearly more robust and appropriate than an artificial segregation of risk analysis steps, implementation of the *Understanding Risk* approach faces both political and practical obstacles.

Sheila Jasanoff is one of the most vocal proponents of broader representation in risk decision making. Her 1994 critique of Breyer's book (Jasanoff 1994), decrying its artificial separation of fact and value in the risk analysis process, points out that most risk decisions are "far too multidimensional to warrant quantification and much too complex to be simulated through any existing computer program." Jasanoff's view is consistent with the recent National Research Council (1996) review of risk analysis philosophy, which argues for eliminating the misleading firewalls between the assessment and management phases. Regardless of where one fits in these debates, a thorough knowledge of current methods is the vital precondition for effective risk analysis.

"Risk comparison," an approach that has been popular in the past decade can be used to exemplify the two risk philosophies. In a technocratic approach, diverse risks are converted to a common metric —perhaps years of life expectancy lost. Risks are ranked along this dimension, and resources are committed to reducing the greatest risks first. Risk comparisons can also, however, be used as a tool to bring together decision makers to discuss how they perceive risks, evaluate the data available to describe those risks, identify the issues upon which they agree and disagree, and decide when decisions can be made and when more information would be useful.

In 1987 and 1990, respectively, the U.S. EPA and its Science Advisory Board used a technocratic approach to review the ways in which environmental risks were prioritized by the existing regulatory legislation and agencies. These studies found that the existing regulations were inconsistent with both expert and lay opinions of the most important risks. Among the reasons for this inconsistency is that environmental regulation evolved piecemeal in response to individual crises, and over several decades. As a result, regulations use disparate approaches for dealing with different media (air, water, foods, facilities). Some statutes call for absolute levels of safety, some require only "prudent" margins, others base standards on current technology, and some require the regulator to balance risks and benefits explicitly. The reports suggested that the EPA's prioritization should be based more explicitly on risk analysis, but absent legislation specifically allowing intermedia risk comparisons, the EPA's options are constrained by existing laws.

Society compares and ranks risks all the time, although often qualitatively and/or implicitly. In a provocative paper, Wilson (1979) asks risk analysts to make some of these comparisons quantitative. Using a one in a million level of risk (where facing a hazard subjects one to a 0.000001 increase in chance of death from that hazard), Wilson compared some everyday and some less common risks. This sort of simple comparison can be eyebrow raising and may usefully question the wisdom of regulating one risk into oblivion at great cost while far larger risks remain unaddressed; however, such point comparisons are limited and highlight the inextricable nature of value judgment.

Table 1-1 indicates that traveling six minutes by canoe is "equal" along this one dimension to living 150 years within twenty miles of a nuclear power plant. But what does, or can, this comparison mean? There is no indication of the certainty associated with the estimates, the (potentially) offsetting benefits, or ways in which they can be avoided. It

Table 1-1 Risks that Increase Chance of Death by 0.000001 (One in One Million, or 10^{-6})

Smoking 1.4 cigarettes	Cancer, heart disease
Drinking 1/2 liter of wine	Cirrhosis of the liver
Spending 1 hour in a coal mine	Black lung disease
Spending 3 hours in a coal mine	Accident
Living 2 days in New York or Boston	Air pollution
Traveling 6 minutes by canoe	Accident
Traveling 10 miles by bicycle	Accident
Traveling 300 miles by car	Accident
Flying 1000 miles by jet	Accident
Flying 6000 miles by jet	Cancer caused by cosmic radiation
Living 2 months in Denver	Cancer caused by cosmic radiation
Living 2 months in average stone or brick building	Cancer caused by natural radiation
One chest X-ray taken in a good hospital	Cancer caused by radiation
Living 2 months with a cigarette smoker	Cancer, heart disease
Eating 40 tablespoons of peanut butter	Liver cancer caused by aflatoxin B
Drinking Miami drinking water for 1 year	Cancer caused by chloroform
Drinking 30 12 oz. cans of diet soda	Cancer caused by saccharin
Living 5 years at site boundary of a typical nuclear power plant in the open	Cancer caused by radiation
Drinking 1000 24 oz. soft drinks from recently banned plastic bottles	Cancer from acrylonitrile monomer
Living 20 years near PVC plant	Cancer caused by vinyl chloride (1976 standard)
Living 150 years within 20 miles of nuclear power plant	Cancer caused by radiation
Eating 100 charcoal broiled steaks	Cancer from benzopyrene
Risk of accident by living within 5 miles of a nuclear reactor for 50 years	Cancer caused by radiation

Source: Wilson 1979.

is often not even clear that such benefits can be calculated. The real insights, and the real work, come from analyses that address the shape and variability of the risk distributions, the confidence associated with each estimate, and the uncertainty generated by data limitations. Until risks are well characterized, it is difficult even to begin comparing.

The environment is by no means the only arena in which risk analysis is receiving increased attention. As the energy sector is deregulated, risk tools have evolved to deal with variability and uncertainty in supply, demand, pricing, and facility design. Similarly, the rash of major catastrophes in the past several years, including seasonal wildfires and major earthquakes in the west, flooding in the midwest, and Hurricane Andrew and beach erosion in the east, has forced the government to incur large expenses, prompted concern about the viability of private insurance underwriting, and promoted more careful risk exposure assessment. As a final example, increased reliance on information technologies has generated concern among public and private decision makers over the security and stability of computer networks. The set of risk analytical tools presented in this volume may be applied to any of these issues.

Cost-benefit analysis (CBA), a version of decision analysis, increasingly accompanies risk analysis on the public policy agenda. Some critics see CBA as nothing more than risk analysis made more complex by adding value judgments such as those putting dollar values on illness, loss of lives, or degradation of ecological resources. (See Costanza et al. 1997 for an example of a truly "grand scale" economic analysis—the value of global natural resources to society—that generates an estimate at the cost of massive uncertainty.) In some cases, simply listing all relevant impacts (positive and negative) without absolute valuations will provide insight into a decision. In other cases, optimizing costs and benefits requires the analyst to quantify all of the tradeoffs in a common metric, usually monetary values. If these choices are, in fact, incommensurate, forcing dollar values on them may be at best arbitrary and at worst self-serving.

Others see CBA as a tool that, while fraught with uncertainty, gives a common rule by which to make necessary comparisons. They argue that society makes these comparisons already, and that CBA will do so in a more consistent, rational manner. The 1994 bill HR9, would have made CBA a legislative requirement, requiring "a final cost-benefit analysis" for every "major rule." Figure 1-1a (after Morgan 1981) represents an idealized economist's approach to solving this problem. However, for many risk issues, the values of the different elements range from extremely difficult to impossible to quantify. How, for example, can we measure the value sports enthusiasts place on the opportunity to play outdoors, and compare this to the costs they impose on society through skin cancer treatment? Figure 1-1b may better represent what we know about many risks.

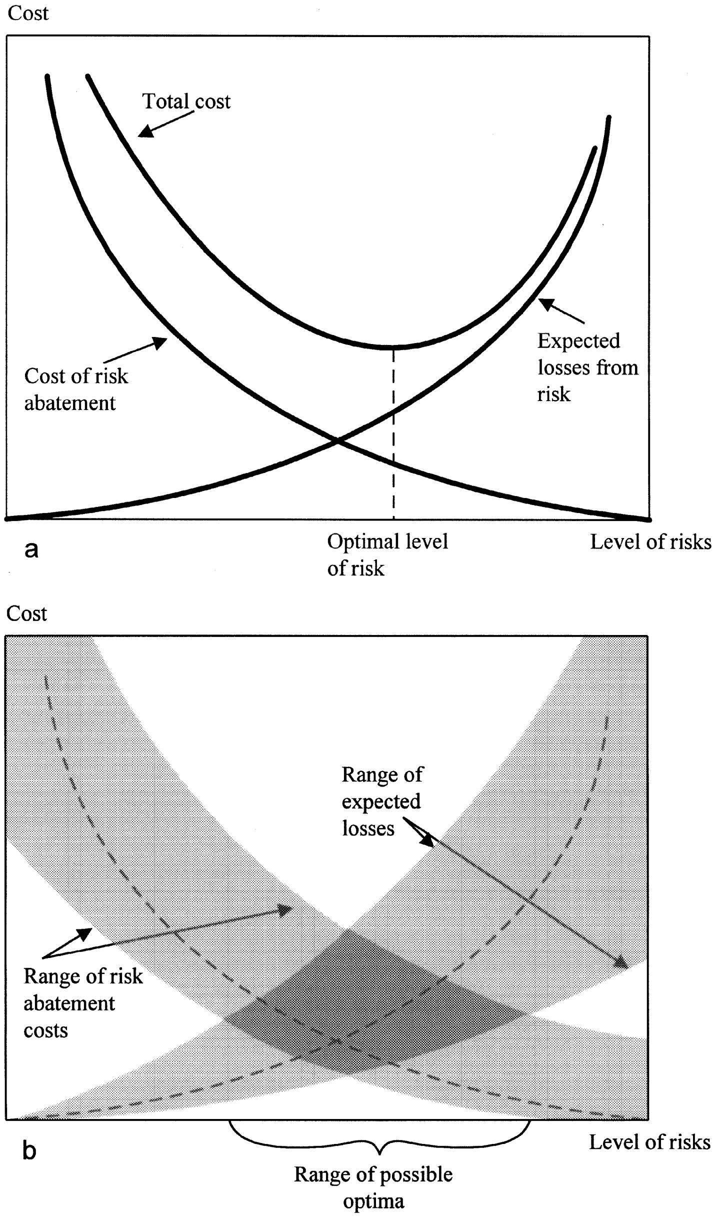

Figure 1-1. (a) depicts how the optimum level of risk and abatement can be calculated, given precise information on costs, benefits, and preferences. (b) suggests that, even if preferences are clearly defined, uncertainties in risk and abatement costs can lead to highly uncertain ranges of possible optima.

In 1985, Professor John Harte of the University of California at Berkeley's Energy and Resources Group created a course on environmental problem solving. Harte's approach was to equip students with a few general tools that allow them to address problems characterized by limited information and apparent complexity. He teaches his students "an approach to problem solving [that] involves the stripping away of unnecessary detail, so that only the essentials remain" (Harte 1985).

Harte presents a three-step approach, a philosophy that he spells out in the preface to *Consider a Spherical Cow* (Harte 1985, pp. xi–xiii). First, he takes a broad overview of a problem (what he calls hand-waving), in order to establish a qualitative understanding of the mechanism of the process being examined. Looking at the "big picture" can often provide an idea of the direction and magnitude of a process, even if the details are obscure. In addition, it can quickly become evident where important information is missing, and which assumptions are most problematic. At this stage, simple "reality checks" suggest whether the solver is on the right track.

Second, he represents the qualitative processes mathematically and uses available data (making assumptions where necessary) to arrive at a "detailed quantitative solution." Third, he evaluates the resilience of his answers if the assumptions he has made are changed or omitted. This step, also called sensitivity analysis, can be applied to both data and assumptions, to suggest where further research will improve understanding and whether uncertainty about the assumptions is likely to overwhelm the results.

In this book, we adapt and extend Harte's environmental problem-solving approach to risk analysis. Harte's philosophy is wonderfully appropriate to risk assessment, where uncertainty is often profound and assumptions must inevitably be made. To familiarize the reader with hand-waving techniques, the following three problems consider important risk problems without using any numbers.

Problem 1-1. Getting Started

Consider figures 1-2 through 1-5. Which of these graphs do you think best represents

a. the number of accidents a driver has, as a function of total cumulative miles driven?

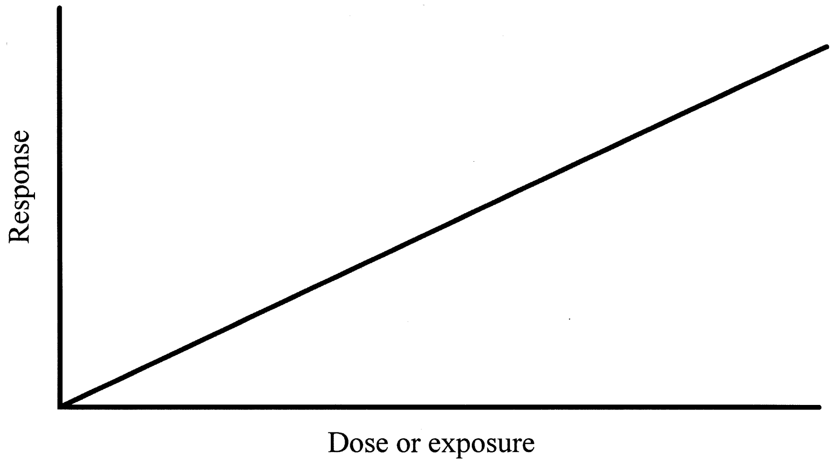

Figure 1-2. A direct relationship between a dose (cause) and its response (effect). It need not be one to one, but it must be true that a unit increase in dose causes a constant increase in response. An example is the purchase of raffle tickets: the chances of winning are directly proportional to the fraction of the total tickets that you hold.

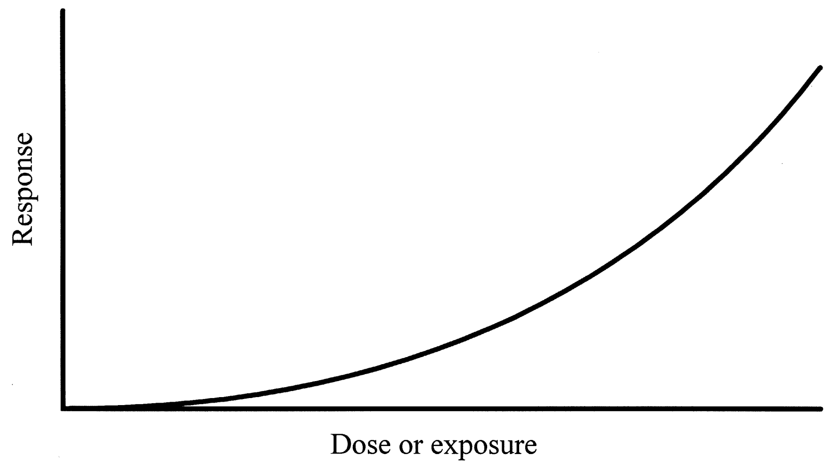

Figure 1-3. A convex relationship, where increases in dose have a relatively larger impact than initial dose. The risk of highway accident per mile traveled as a function of travel speed is a convex relationship. Convex functions are those for which the second derivative is positive, meaning that the slope of the line increases throughout the convex region.

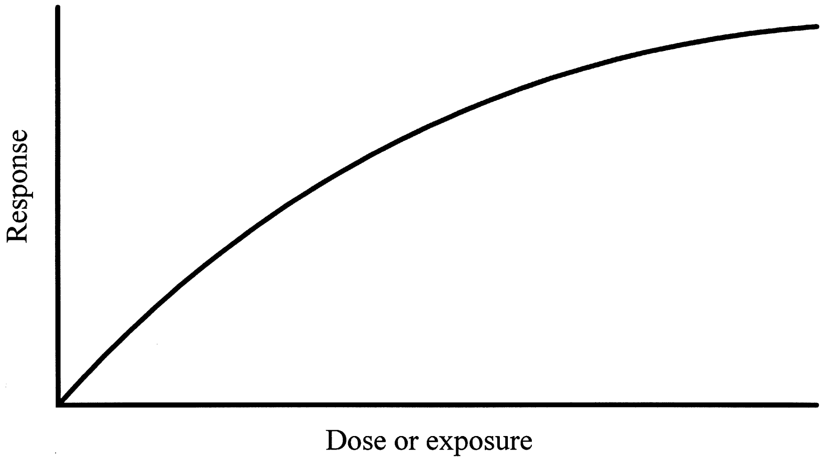

Figure 1-4. A concave relationship where additional dose has a smaller relative response than does the initial dose. This curve describes, for example, the relationship between pedaling effort and bicycling speed: wind resistance increases at a faster rate than the increase in speed. Concave functions are those for which the second derivative is negative, meaning that the slope of the line decreases throughout the concave region.

b. the number of space shuttle accidents as a function of total number of missions flown?
c. the number of leaks in a sewer line as a function of the number of years it has been in service without maintenance or replacement?
d. the number of carcinomas a surfer is likely to get as a function of total lifetime hours in the sun?

Solution 1-1

Differing interpretations are possible for several of the cases, and without further analysis we are, at this point, making educated guesses. However, there can often be considerable value to these educated or hand-waving guesses. It is important not to wave your hands frantically like a lost hiker hoping to be spotted by a passing plane, but rather in the controlled and directed fashion of a symphony conductor. Thinking in very general terms can often point out where we have good enough information to make a reasonable decision, and where our guesses are so broad or unrealistic that we must push the questions further.

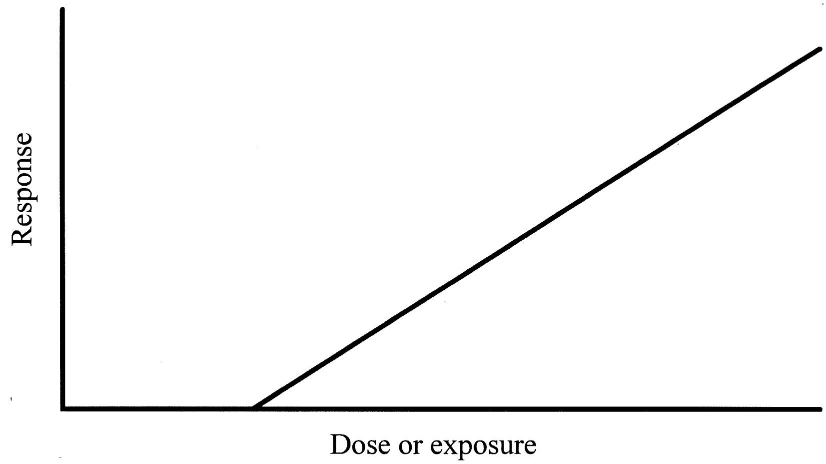

Figure 1-5. A threshold relationship in which initial doses have no effect, but eventually a dose (the threshold) is reached beyond which a response is elicited. An example is a redundant system, where the first several failures are protected by backups, but eventually the backups are exhausted and effects begin to appear. Note that threshold effects can include linear, convex, or concave patterns.

Solution 1-1a

Since drivers tend to gain experience over time, figure 1-3 (concave) is probably a good representation. The driver's total number of accidents will grow, but at a decreasing rate. In this context, the concave response is often referred to as a "learning curve," as the failures grow less frequent with experience, and there are diminishing marginal benefits from experience. In some physical contexts this curve describes saturation, where subsequent doses do not elicit as much response.

Solution 1-1b

Assuming that routine maintenance is done, and the launch- and space-worthiness maintained, figure 1-2 (linear) might be a good choice (see point one below). There is a roughly constant probability of an accident for each launch, which leads to a linear cumulative hazard.

Solution 1-1c

One would expect a new sewer line to have a fairly high integrity, which would suggest a very small number of leaks early on. However, as it continues in service without maintenance, ground shifting, corrosion, and other effects cause increasing numbers of leaks, and probably at an increasing rate. Consequently, figure 1-4 (convex), an exponentially increasing number of leaks over time, may be a good model, or possibly figure 1-5 (threshold), if it takes a while for the first leak to start, and then the effects of the corrosion from the initial leaks exacerbate the effects of other deleterious forces.

Solution 1-1d

There are arguments for any of the four models here; the same is true for many causes of cancer and other types of chemical toxicity. Figure 1-2 makes sense if each unit of energy is equally likely to create a cancer cell. Figure 1-3 would apply if there were some saturation effect where most of the damage is due to the initial exposure to radiation. Figure 1-4 implies that additional radiation is likely to exacerbate the effect of earlier exposure, meaning an increasing rate of carcinomas over time. Finally, figure 1-5 (threshold) represents the case where the body is able to repair the damage due to a limited amount of radiation (or, for example, to metabolize a toxin up to some amount), but beyond that threshold level, it cannot, and carcinogenesis begins. The question of whether a process has a threshold is fundamental to quantitative and qualitative risk analysis.

Problem 1-1 highlights several issues. First, there is often insufficient information to come up with an indisputably "right" answer. We make a number of assumptions about what is causing the phenomena we are interested in and use these assumptions to theorize about the expected outcome. Depending on what aspects we think are most important and relevant, we may individually arrive at quite different sets of assumptions. For example, you may have assumed that there is a learning curve associated with space shuttle launches, in which case the curve would be concave. On the other hand, if you thought that the individual shuttles were likely to be subject to wear and tear, you might predict an exponential increase in the number of accidents.

Second, based on our theories about how things work, we "build models." Throughout this book, you should keep in mind that the concept of modeling is simple: use what we know to describe what we observe. The models in this exercise were built without using equations or numbers, but we have an idea of what is on each axis and how the axes relate. Plugging in the numbers (provided we can get them) and calculating can often be a trivial exercise; the important point is to understand what is going on. The benefit of eventually plugging in numbers is that we can use them to predict future outcomes. While this ability to predict is the goal of risk assessment, differences in assumptions and theories can lead to highly divergent numbers, that is to say, uncertainty.

Third, a single model can be used to describe very different phenomena. The essential modeling relationships of, for example, carcinogenesis and automobile accidents may be analogous, even when the physical processes are entirely different. This is extremely important, because it allows us to develop general methods for thinking about a wide range of problems. The next step is to refine the models and make them better fit the specific case under scrutiny. This in turn requires more information.

Problem 1-2. Data Needs

What evidence would you want to confirm your (or our) answers to problem 1-1?

Solution 1-2

Problem 1-1 is about constructing theoretical models; problem 1-2 is about verifying and calibrating the models empirically by comparing them to data. Two things to keep in mind are how well the data fit the model and how "good" the data are . . . and recall that bad data may erroneously "confirm" a bad theory!

Solution 1-2a

The insurance industry has an abundance of data on this subject. In general, younger drivers tend to get into more accidents than do older drivers, but the decreasing trend tends to plateau at some age over thirty. Note that there may be two mechanisms operating here. One is experience—the number of years that an individual has been driving, and the skills he or she has gained through that practice. The other is maturity—older drivers may be less risk-taking relative to younger drivers. Note also that insurance companies usually use individual data as well, such as number and magnitude of prior accidents.

Solution 1-2b

Richard Feynman, in a 1988 article documenting his review of the space shuttle Challenger 1985 explosion, found that some of the engineers estimated about a one in two hundred chance of such a failure, based on their understanding of the materials and very complex equipment involved. Meanwhile, people at higher levels in the administration assumed much smaller probabilities, on the order of one in ten thousand. The accident occurred on the seventy-eighth flight, and while limited inference is possible given only one occurrence, the engineers' model appears better supported than that of the administration. Why the difference? It is likely—as the subsequent investigation showed—that the politics and finances of the shuttle program exerted a strong pressure to remain "on schedule." Thus, while more information would improve decision making about this particular risk, we are not likely to get it in time to make good decisions. (In fact, while additional safe flights extend the data set, we hope we do not get additional "failure" data!) Consequently, the choice of the right model will be based on the extent to which we believe the assumptions behind each. The moral of the story is that not all assumptions can be tested. A large number of failures makes modeling easier, while few (or no) failures makes predictions extremely uncertain.

Solution 1-2c

Data on this could come from a variety of sources. Ideally, one would want to inspect the pipe in question regularly and check for leaks. Alternatively, a comparison to similar pipes in similar use might provide relevant information. Laboratory tests on the pipe may give information about the susceptibility to failure over time, and geologic history might suggest the types of stresses the pipe is subject to. Manufacturers and sewer companies may have historical and laboratory data on material specifications.

Solution 1-2d

The types of data needed to support one answer or another come in two broad classes: toxicological and epidemiological. We will consider both of these in much greater detail later in the book. In general, a toxicological test would involve exposing groups of individuals to varying levels of sunlight (or ultraviolet light), while keeping everything else in their lives the same, to see whether different skin cancer levels result. Since it is difficult (and ethically unacceptable) to do this sort of test on people, it is more often done on small groups of animals or cultured human cells. Additional assumptions must therefore be made, for example, about the relationship between animal and human carcinogen susceptibility.

An epidemiological test would try to find individuals who have been exposed to different levels of sunlight in the past, and compare rates of skin cancer among those groups. The problems here are likely to be with data quality, such as accuracy in determining past exposure, and in insuring that some third factor that was not measured is not the cause.

Problem 1-3. Using Data

Assume that you found that for problem 1-1d, figure 1-3 was the most likely model, and that 50% of people who surfed regularly could be expected to get at least one carcinoma. How would this affect your attitude about surfing? How would you tell people about your findings? Would you suggest that surfing be regulated?

Solution 1-3

First, note that this issue will come up again much later in the book, when evaluating the relationships between numbers, such as one in a million, 1%, 10%, 50%, and more abstract concepts, such as “rare,” "very rare," and "common," and what these distinctions and relationships mean when it comes to making policy.

Clearly, there is no obvious answer to this question. However, where the first two problems asked us to formalize our thinking about risk, this one asks what to do once we have information. We need to think about how people perceive risks—does the average surfer think she has a 50/50 chance of getting cancer? We need to think about how to advise surfers about their risks. We also need to think about how and whether to try to compare this cancer risk to other risks and benefits, how and whether surfing fits into our social and regulatory system, and whether and what additional information would improve our decisions. Finally, we need to think about which surfers are at risk. Does each of them have a 50/50 chance to contract cancer, or are some of the surfers at higher risk for genetic or other reasons? If people do vary, can we figure out which are the high-risk surfers? And if we can determine the high-risk individuals, can and should we treat them differently? Finally, what actions can surfers take?—sunscreen, wet suits, T-shirts, and so on.

Problem 1-A. Additional Cases

Answer problems 1-1 and 1-2 for the following cases:

a. The expected number of "heads" from tosses of a fair coin as a function of total number of tosses
b. The probability of a ski jumper crashing as a function of the height of the mogul from which she jumps
c. The number of times a cheap handgun will misfire as a function of the number of times it is fired

---

Problem 1-B. Additional Curves

a. Consider the observation that at some stage in their lives, the competence of many drivers begins to deteriorate. Draw a curve that represents a driver's lifetime driving experience, beginning with a steep but gradually leveling learning curve, followed by a long period of no change, followed by increasing risk late in life.

b. For some processes and items—washing machines, for example—there is some chance that the item will fail upon initial use. However, if the item operates successfully the first time, the additional chance of failure grows very slowly over time. Graph this phenomenon.

c. Suggest some additional curves, along with cases that they might represent.

Problem 1-C. Does the Dose Make the Poison?

a. Assume that newborn body weight is a reasonable measure of health, with higher weight meaning better health, and that vitamin D is essential in moderate quantities but injurious in excess; implying an optimum dose corresponding to a maximum average birth weight. Propose a curve that would represent newborn body weight as a function of the mother's intake of vitamin D.

b. Repeat (a) using reduction in body weight instead of body weight. How do the curves differ? In what respects are they the same?

c. Rephrase (a) such that it asks the same question but is represented by a (seemingly) different curve. Discuss the effects the different representations might have on perceptions of the effects.

Problem 1-D. One in a Million Risks

Refer back to table 1-1.

a. Why are there two different entries for risk from coal mines? How would this table seem different if the two were combined into a single 10^{-6} risk of spending three quarters of an hour in a coal mine?

b. According to this table, if you lived in Miami and drank a can of diet soda each week, your risk from that would be greater than that from drinking the tap water. Does this mean that you are unreasonable if you object to carcinogens in your drinking water?

That you are taking on too much risk from diet soda? Neither? Discuss.

Problem 1-E. Surfing and Smoking

The debate about managing cancer and other risks associated with cigarettes is both similar to and dissimilar from that about skin cancer and other risks from surfing. Compare the question of public health and surfing to that of public health and smoking. In what ways do the issues differ, and in what ways are they the same?

Problem 1-F. Risks of Nuclear Power

The operation of large (1-GW_e, or 10^a watts, scale) nuclear power reactors began in about 1970 and there are approximately 350 nuclear power reactors operating worldwide today.

a. Roughly how many nuclear power reactor-years of operation have accumulated during this period?
b. During this period, there has been one accident that resulted in a major release of radioactivity (Chernobyl 1986) and one accident in which all but a small amount of release was prevented by the reactor containment building (Three-Mile Island 1979). On this basis, roughly what is the probability of a major release per reactor-year?
c. In 1975, the WASH-1400 Reactor Safety Study Report (see box 1-2 and figures 1-6 and 1-7) estimated that the chances of a Chernobyl-type accident were around one in one million each year. Is this consistent with your estimate in (b)? What might account for the differences?
d. The Chernobyl accident may ultimately cause on the order of 10,000 extra cancer deaths (von Hippel and Cochrane 1991). How many would this come to per reactor-year?

Box 1-2. The WASH 1400 report

In 1975, the United States Nuclear Regulatory Commission completed a now famous (or infamous) report entitled "Reactor Safety Study," which tried to quantify the probability of various types of reactor accidents that might occur, and compare those risks to other risks that people already face. The report continues to be discussed because it generated both great applause and criticism—sometimes both from the same individuals! The report's proponents argue that it was the most thorough and quantitatively rigorous risk analysis ever done in any context, and provided the best possible numbers to inform policy makers. Its detractors counter that it completely ignored potentially disastrous interactions, and used comparisons (such as tornadoes versus core meltdowns) that ignored inherent differences.

In addition, while including immediate radiation-induced deaths caused by various possible failure scenarios, figures 1-6 and 1-7 failed to depict associated long term cancer deaths. As in the Chernobyl case, the latter could be orders of magnitude greater than the number of short-term deaths (von Hippel and Cochran 1991). Consequently, the curves on figures 1-6 and 1-7 are highly misleading. For example, a curve taking into account long-term cancer deaths would predict about 10,000 deaths at a frequency of 10^{-3} per year for 100 reactors, rather than the figure 1-6 and 1-7 predictions, which approaches zero fatalities at that frequency! (Hohenemser et al. 1992).

e. Each year in the United States there are roughly fifty fatalities in coal mines, a few hundred in coal transport, and a few thousand due to respiratory diseases caused by the emission of SO_2 by the equivalent of three hundred 1-GW_e U.S. coal-fired power plants. On this (quite uncertain) basis, approximately how many coal-fired power plant fatalities are there per equivalent reactor-year?

f. Society does not respond in the same way to all risks, for example, the risks from nuclear and coal power plants—or to the fatalities from auto accidents, smoking, and skiing. Discuss why this appears to be the case.

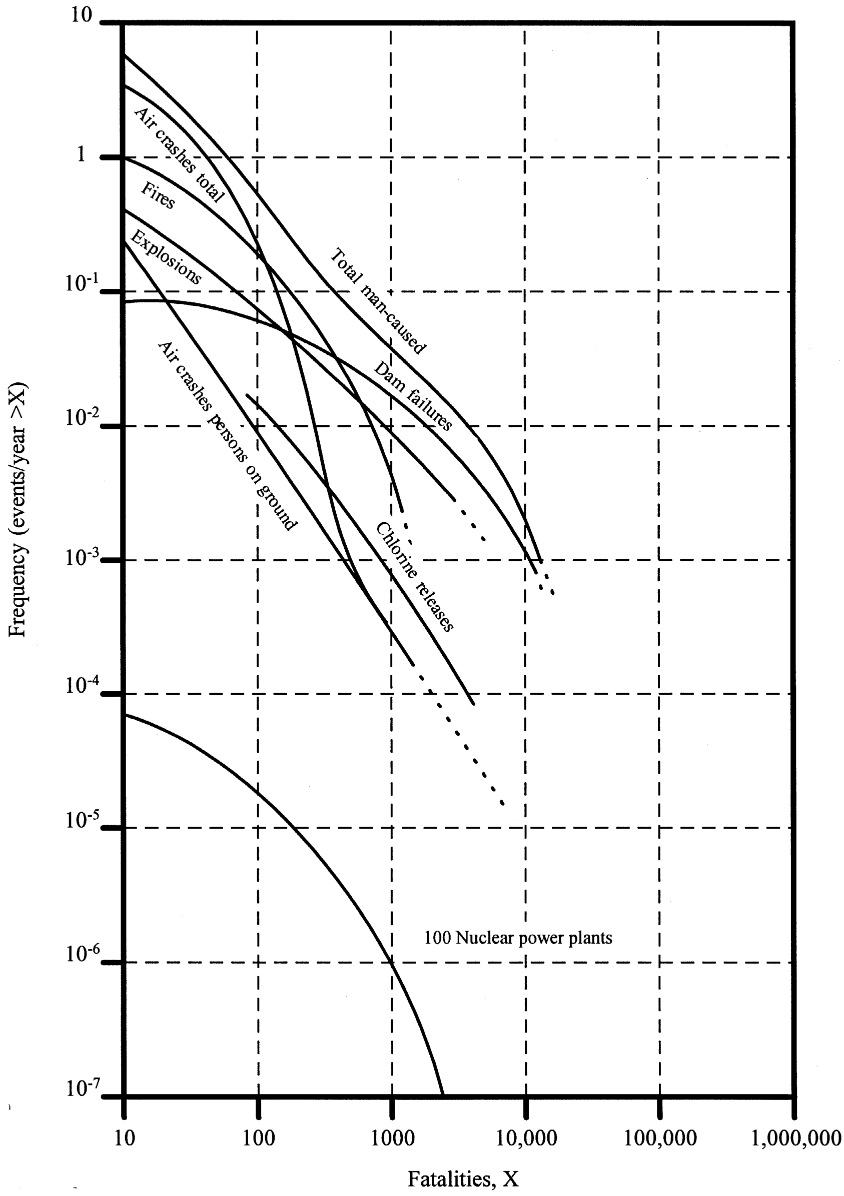

Figure 1-6. Frequency of man-caused events involving fatalities. Figure 6-1 from WASH-1400 report (U.S. Nuclear Regulatory Commission 1975).

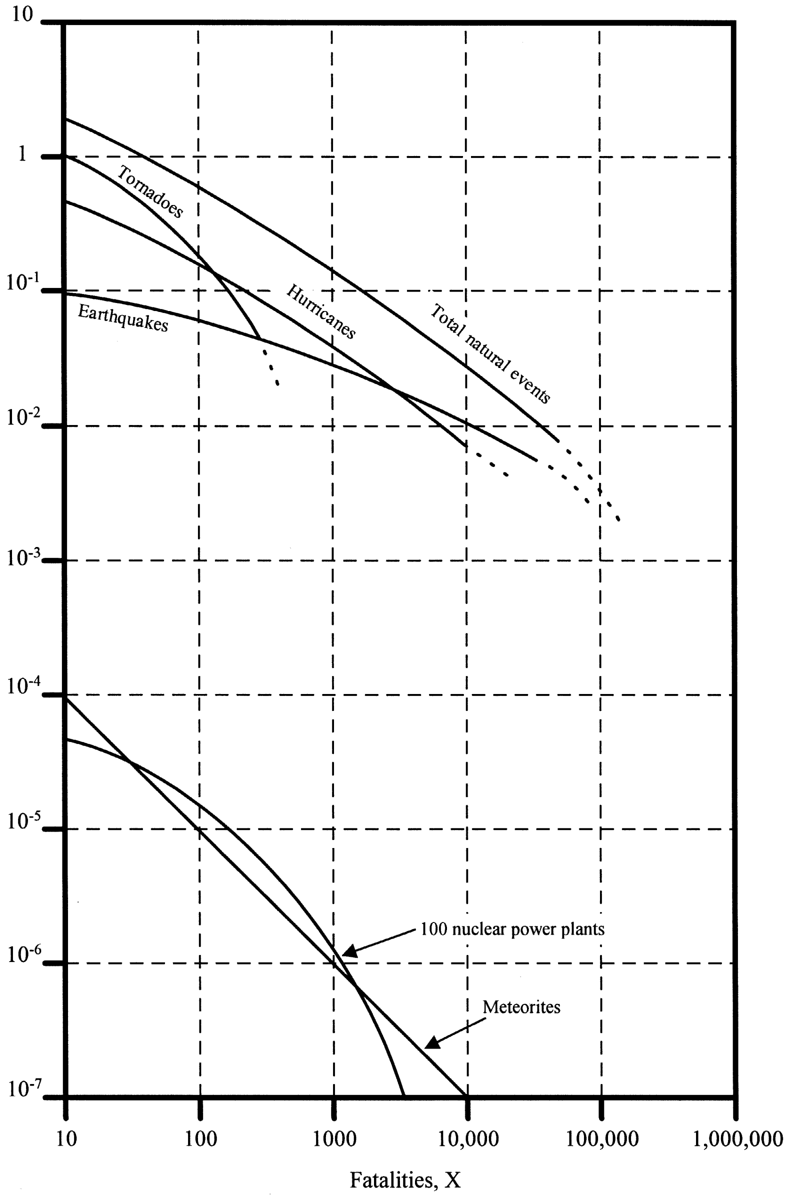

Figure 1-7. Frequency of natural events involving fatalities. Figure 6-2 from WASH-1400 report (from U.S. Nuclear Regulatory Commission 1975).

References

Breyer, S. G. (1993). *Breaking the Vicious Circle: Toward Effective Risk Regulation*. Cambridge, MA: Harvard University Press.

Costanza R., d'Arge, R., de Groot, R., Farber, S., Grasso, M., Hannon, B., Limburg, K., Naeem, S., O'Neill, R., Paruelo, J., Raskin, R, Sutton, P., and van den Belt, M. (1997), "The value of the world's ecosystem services and natural capital." *Nature* 387:253–60.

Covello, V. T. and Mumpower, J. (1985). "Risk analysis and risk management: An historical perspective," *Risk Analysis* 5(2):103–20.

Feynmann, R. P. (1988), "An outsider's inside view of the Challenger inquiry," *Physics Today* (February) 26–37.

Harte, J. (1985). *Consider a Spherical Cow: a Course in Environmental Problem Solving*. Mill Valley, CA: University Science Books.

Hohenemser, C., Goble, R., and Slovic, P. (1992). "Nuclear power." In Hollander, J. M. (Ed.) *The Energy-Environment Connection*. Covello, CA: Island Press.

Jasanoff, S. (1993). "Bridging the two cultures of risk analysis." *Risk Analysis* 13(2):123–29.

———. (1994). "The dilemmas of risk regulation." *Issues in Science and Technology Policy* (Spring):79–81.

Morgan, M. G. (1981). "Choosing and managing technology-induced risk," *IEEE Spectrum* 18(12):53–60.

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

———. (1994). *Science and Judgment in Risk Assessment*. Washington, DC: National Academy Press.

———. (1996). *Understanding Risk*. Washington, DC: National Academy Press.

Otway, H. (1992). "Public wisdom, expert fallibility." In Krimsky, S. and Golding, D. (Eds.) *Social Theories of Risk*. NY: Praeger.

U.S. EPA (1987). *Unfinished Business: A Comparative Assessment of Environmental Problems*. Washington, DC: U.S. EPA Office of Policy Analysis.

———. (1990). *Reducing Risk: Setting Priorities and Strategies for Environmental Protection*. Washington, DC: U.S. EPA Science Advisory Board.

U.S. Nuclear Regulatory Commission. (1975). *Reactor Safety Study: An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants*. Executive Summary, WASH-1400, NUREG-75/014.

von Hippel, F. N. and Cochran, T. (1991). "Chernobyl: estimating the long term health effects." In *Citizen Scientist* New York, NY: Touchstone Books.

Weinberg, Alvin. (1972). "Science and trans-science." *Minerva* 10(2):209–21.

Wilson, R. (1979). "Analyzing the daily risks of life." *Technology Review* 81(4):41–46.