

Risk Analysis and Management

Inadequate approaches to handling risks may result in bad policy. Fortunately, rational techniques for assessment now exist

M. Granger Morgan

July 1993

Contents

Risk Analysis in Action 14
Further Reading 16

M. GRANGER MORGAN has worked for many years to improve techniques for analyzing and managing risks to health, safety and the environment. Morgan heads the department of engineering and public policy at Carnegie Mellon University. He also holds appointments in the department of electrical and computer engineering and at the H. John Heinz III School of Public Policy and Management. Morgan received a B.A. from Harvard University, an M.S. from Cornell University and a Ph.D. in applied physics from the University of California, San Diego.



AIR DISASTER in Madrid claimed 183 lives in November 1983. The (small) chance of dying in an air crash is one of the prices that society agrees to pay for rapid, convenient global transportation. Some risks, including nuclear power generation, have caused fewer deaths but provoked greater calls for regulation, whereas others, such as automobiles, cause more deaths but arouse less concern.

Americans live longer and healthier lives today than at any time in their history. Yet they seem preoccupied with risks to health, safety and the environment. Many advocates, such as industry representatives promoting unpopular technology or Environmental Protection Agency staffers defending its regulatory agenda, argue that the public has a bad sense of perspective. Americans, they say, demand that enormous efforts be directed at small but scary-sounding risks while virtually ignoring larger, more commonplace ones.

Other evidence, however, suggests that citizens are eminently sensible about risks they face. Recent decades have witnessed precipitous drops in the rate and social acceptability of smoking, widespread shifts toward low-fat, high-fiber diets, dramatic improvements in automobile safety and the passage of mandatory seat belt laws—all steps that reduce the chance of untimely demise at little cost.

My experience and that of my colleagues indicate that the public can be very sensible about risk when companies, regulators and other institutions give it the opportunity. Laypeople have different, broader definitions of risk, which in important respects can be more rational than the narrow ones used by experts. Furthermore, risk management is, fundamentally, a question of values. In a democratic society, there is no acceptable way to make these choices without involving the citizens who will be affected by them.

The public agenda is already crowded with unresolved issues of certain or potential hazards such as AIDS, asbestos in schools and contaminants in food and drinking water. Meanwhile scientific and social developments are bringing new problems—global warming, genetic engineering and others—to the fore. To meet the challenge that these issues pose, risk analysts and managers will have to change their agenda for evaluating dangers to the general welfare; they will also have to adopt new communication styles and learn from the populace rather than simply trying to force information on it. While public trust in risk management has declined, ironically the discipline of risk analysis has matured. It is now possible to examine potential hazards in a rigorous, quantitative fashion and thus to give people and their representatives facts on which to base essential personal and political decisions.


Risk analysts start by dividing hazards into two parts: exposure and effect. Exposure studies look at the ways in which a person (or, say, an ecosystem or a piece of art) might be subjected to change; effects studies examine what may happen once that exposure has manifested itself. Investigating the risks of lead for inner-city children, for example, might start with exposure studies to learn how old, flaking house paint releases lead into the environment and how children build up the substance in their bodies by inhaling dust or ingesting dirt. Effects studies might then attempt to determine the reduction in academic performance attributable to specific amounts of lead in the blood.

Exposure to a pollutant or other hazard may cause a complex chain of events leading to one of a number of effects, but analysts have found that the overall result can be modeled by a function that assigns a single number to any given exposure level. A simple, linear relation, for instance, accurately describes the average cancer risk incurred by smokers: 10 cigarettes a day generally increase the chance of contracting lung cancer by a factor of 25; 20 cigarettes a day increase it by a factor of 50. For other risks, however, a simple dose-response function is not appropriate, and more complex models must be used.

The study of exposure and effects is fraught with uncertainty. Indeed, uncertainty is at the heart of the definition of risk. In many cases, the risk may be well understood in a statistical sense but still be uncertain at the level of individual events. Insurance companies cannot predict whether any single driver will be killed or injured in an

accident, even though they can estimate the annual number of crash-related deaths and injuries in the U.S. with considerable precision.


For other risks, such as those involving new technologies or those in which bad outcomes occur only rarely, uncertainty enters the calculations at a higher level—overall probabilities as well as individual events are unpredictable. If good actuarial data are not available, analysts must find other methods to estimate the likelihood of exposure and subsequent effects. The development of risk assessment during the past two decades has been in large part the story of finding ways to determine the extent of risks that have little precedent.




In one common technique, failure mode and effect analysis, workers try to identify all the events that might help cause a system to break down. Then they compile as complete a description as possible of the routes by which those events could lead to a failure (for instance, a chemical tank might release its contents either because a weld cracks and the tank ruptures or because an electrical short causes the cooling system to stop, allowing the contents to overheat and eventually explode). Although enumerating all possible routes to failure may sound like a simple task, it is difficult to exhaust all the alternatives. Usually a system must be described several times in different ways before analysts are confident that they have grasped its intricacies, and even then it is often impossible to be sure that all avenues have been identified.

Once the failure modes have been enumerated, a fault tree can aid in estimating the likelihood of any given mode. This tree graphically depicts how the subsystems of an object depend on one another and how the failure of one part affects key operations. Once the fault tree has been constructed, one need only estimate the probability that individual elements will fail to find the chance that the entire system will cease to function under a particular set of circumstances. Norman C. Rasmussen of the Massachusetts Institute of Technology was among the first to use the method on a large scale when he directed a study of nuclear reactor safety in 1975. Although specific

EXPOSURE PROCESSES



EFFECTS PROCESSES




RISK MANAGEMENT PROCESS begins with analysis of the people and other entities exposed to change, such as in this illustration, from emissions from a coal-burning power plant (*left*). After the results of exposure have been quantified (*second panel*), they must then be filtered through public perceptions, which cause people to respond more strongly to some aspects of risk than to others. Ultimately, costs and benefits will be weighed. Agreeing on the values used to make decisions and making sure that all relevant effects are taken into account are crucial, but often neglected, parts of the process.

details of his estimates were disputed, fault trees are now used routinely in the nuclear industry and other fields.

Boeing applies fault-tree analysis to the design of large aircraft. Company engineers have identified and remedied a number of potential problems, such as vulnerabilities caused by routing multiple control lines through the same area. Alcoa workers recently used fault trees to examine the safety of their large furnaces. On the basis of their findings, the company revised its safety standards to mandate the use of programmable logic controllers for safety-critical controls. They also instituted rigorous testing of automatic shut-off valves for leaks and added alarms that warn operators to close manual isolation valves during shutdown periods. The company estimates that these changes have reduced the likelihood of explosions by a factor of 20. Major chemical companies such as Du Pont, Monsanto and Union Carbide have also employed the technique in designing processes for chemical plants, in deciding where to build plants and in evaluating the risks of transporting chemicals.

In addition to dealing with uncertainty about the likelihood of an event such as the breakdown of a crucial piece of equipment, risk analysts must cope with other unknowns: if a chemical tank leaks, one cannot determine beforehand the exact amount of pollutant released, the precise shape of the resulting dose-response curves for people exposed, or the values of the rate constants governing the chemical reactions that convert the contents of the tank to more or less dangerous forms. Such uncertainties are often represented by means of probability distributions, which describe the odds that a quantity will take on a specific value within a range of possible levels.

When risk specialists must estimate the likelihood that a part will fail or assign a range of uncertainty to an essential value in a model, they can sometimes use data collected from similar systems elsewhere—although the design of a proposed chemical plant as a whole may be new, the components in its high-pressure steam systems will basically be indistinguishable from those in other plants.



SUPERCOMPUTER MODEL of ozone concentrations in the Los Angeles basin (*pink, highest; yellow, lowest*) serves as a starting point for analyses of the risks of exposure to air pollutants.

In other cases, however, historical data are not available. Sometimes workers can build predictive models to estimate probabilities based on what is known about roughly similar systems, but often they must rely on expert subjective judgment. Because of the way people think about uncertainty, this approach may involve serious biases. Even so, quantitative risk analysis retains the advantage that judgments can be incorporated in a way that makes assumptions and biases explicit.

Only a few years ago such detailed study of risks required months of custom programming and days or weeks of mainframe computer time. Today a variety of powerful, general-purpose tools are available to make calculations involving uncertainty. These programs, many of which run on personal computers, are revolutionizing the field. They enable accomplished analysts to complete projects that just a decade ago were considered beyond the reach of all but the most sophisticated organizations [*see box*

on page 38]. Although using such software requires training, they could democratize risk assessment and make rigorous determinations far more widely available.

After they have determined the likelihood that a system could expose Y people to harm and described the particulars of the damage that could result from exposure, some risk analysts believe their job is almost done. In fact, they have just completed the preliminaries. Once a risk has been identified and analyzed, psychological and social processes of perception and valuation come into play. How people view and evaluate particular risks determines which of the many changes that may occur in the world they choose to notice and perhaps do something about. Someone must then establish the rules for weighing risks, for deciding if the risk is to be controlled and, if so, how. Risk management thus tends to force a society to consider what it cares about and who should bear the burden of living with or mitigating a problem once it has been identified.

For many years, most economists and technologists perceived risk simply in terms of expected value. Working for a few hours in a coal mine, eating peanut butter sandwiches every day for a month, and living next to a nuclear power plant for five years all involve an increased risk of death of about one in a million, so analysts viewed them all as equally risky. When people are asked to rank various activities and technologies in terms of risk, however, they produce lists whose order does not correspond very closely to the number of expected deaths. As a result, some early risk analysts decided that people were confused and that their opinions should be discounted.

Since then, social scientists have conducted extensive studies of public risk perception and discovered that the situation is considerably more subtle. When people are asked to order well-known hazards in terms of the number of deaths and injuries they cause every year, on average they can do it pretty well. If, however, they are asked to rank those hazards in terms of risk, they produce quite a different order.

People do not define risk solely as the expected number of deaths or injuries per unit time. Experimental psychologists Baruch Fischhoff of Carnegie Mellon University and Paul Slovic and Sarah Lichtenstein of Decision Research in Eugene, Ore., have shown that people also rank risks based on how well the process in question is understood, how equitably the danger is distributed, how well individuals can control their exposure and whether risk is assumed voluntarily.

Slovic and his colleagues have found that these factors can be combined into three major groups. The first is basically an event's degree of dreadfulness (as determined by such features as the scale of its effects and the degree to which it affects "innocent" bystanders). The second is a measure of how well the risk is understood, and the third is the number of people exposed. These groups of characteristics can be used to define a "risk space." Where a hazard falls within this space says quite a lot about how people are likely to respond to it. Risks carrying a high level of "dread," for example, provoke more calls for government intervention than do some more workaday risks that actually cause more deaths or injuries.

In making judgments about uncertainty, including ones about risk, experimental psychologists have found that people unconsciously use a number of heuristics. Usually these rules of thumb work well, but under some circumstances they can lead to systematic bias or other errors. As a result, people tend to underestimate the frequency of very common causes of death—stroke, cancer, accidents—by roughly a factor of 10. They also overestimate the frequency of very uncommon causes of death (botulism poisoning, for example) by as much as several orders of magnitude.

These mistakes apparently result from the so-called heuristic of availability. Daniel Kahneman of the University of California at Berkeley, Amos N. Tversky of Stanford University and others have found that people often judge the likelihood of an event in terms of how easily they can recall (or imagine) examples. In this case, stroke is a very common cause of death, but most people learn about it only when a close friend or relative or famous person dies; in contrast, virtually every time someone dies of botulism, people are likely to hear about it on the evening news. This heuristic and others are not limited to the general public. Even experts sometimes employ them in making judgments about uncertainty.

Once people have noticed a risk and decided that they care enough to do something about it, just what should they do? How should they decide the amount to be spent on reducing the risk, and on whom should they place the primary burdens? Risk managers can intervene at many points: they can work to prevent the process producing the risk, to reduce exposures, to modify effects, to alter perceptions or valuations through education and public relations or to compensate for damage after the fact. Which strategy is best depends in large part on the attributes of the particular risk.

Even before determining how to intervene, risk managers must choose the rules that will be used to judge whether to deal with a particular issue and, if so, how much attention, effort and money to devote. Most rules fall into one of three broad classes: utility based, rights based and technology based. The first kind of rules attempt to maximize net benefits. Analysts add up the pros and cons of a particular course of action and take the difference between the two. The course with the best score wins.

Early benefit-cost analyses employed fixed estimates of the value of good and bad outcomes. Many workers now use probabilistic estimates instead to reflect the inherent uncertainty of their descriptions. Although decisions are ultimately made in terms of expected values, other measures may be employed as well. For example, if the principal concern is to avoid disasters, analysts could adopt a “minimax” criterion, which seeks to minimize the harm done by the worst possible outcome, even if that leads to worse results on average.

Of course, many tricky points are involved in such calculations. Costs and benefits may not depend linearly on the amount of pollutant emitted or on the number of dollars spent for control. Furthermore, not all the pros and cons of an issue can necessarily be measured on the same scale. When the absolute magnitude of net benefits cannot be estimated, however, rules based on relative criteria such as cost-effectiveness can still aid decision makers.

Rights-based rules replace the notion of utility with one of justice. In most utility-based systems, anything can be subject to trade-offs; in rights-based ones, however, there are certain things that one party cannot do to another without its consent, regardless of costs or benefits. This is the approach that Congress has taken (at least formally) in the Clean Air Act of 1970: the law does not call for maximizing net social benefit; instead it just requires controlling pollutant concentrations so as to protect the most sensitive populations exposed to them. The underlying presumption holds that these individuals have a tight to protection from harm.

Technology-based criteria, in contrast to the first two types, are not concerned with costs, benefits or rights but rather with the level of technology available to control certain risks. Regulations based on these criteria typically mandate “the best available technology” or emissions that are “as low as reasonably achievable.” Such rules can be difficult to apply because people seldom agree on the definitions of “available” or “reasonably achievable.” Furthermore, technological advances may impose an unintended moving target on both regulators and industry.

There is no correct choice among the various criteria for making decisions about risks. They depend on the ethical and value preferences of individuals and society at large. It is, however, critically important that decision frameworks be carefully and explicitly chosen and that these choices be kept logically consistent, especially in complex situations. To do otherwise may produce inconsistent approaches to the same risk. The epa has slipped into this error by writing different rules to govern exposure to sources of radioactivity that pose essentially similar risks.

Implicit in the process of risk analysis and management is the crucial role of communication. If public bodies are to make good decisions about regulating potential hazards, citizens must be well informed. The alternative of entrusting policy to panels of experts working behind closed doors has proved a failure, both because the resulting policy may ignore important social considerations and because it may prove impossible to implement in the face of grass-roots resistance.

Until the mid-1980s, there was little research on communicating risks to the public. Over the past five years, along with my colleagues Fischhoff and Lester B. Lave, I have found that much of the conventional wisdom in this area does not hold up. The chemical industry, for example, distilled years of literature about communication into advice for plant managers on ways to make public comparisons between different kinds of risks. We subjected the advice to empirical evaluation and found that it is wrong. We have concluded that the only way to communicate risks reliably is to start by learning what people already know and what they need to know, then develop messages, test them and refine them until surveys demonstrate that the messages have conveyed the intended information.

In 1989 we looked at the effects of the epa’s general brochure about radon in homes. The epa prepared this brochure according to traditional methods: ask scientific experts what they think people should be told and then package the result in an attractive form. In fact, people are rarely completely ignorant about a risk, and so they filter any

message through their existing knowledge. A message that does not take this filtering process into account can be ignored or misinterpreted.

To study people's mental models, we began with a set of open-ended interviews, first asking, "Tell me about radon." Our questions grew more specific only in the later stages of the interview. The number of new ideas encountered in such interviews approached an asymptotic limit after a couple of dozen people. At this point, we devised a closed-form questionnaire from the results of the interviews and administered it to a much larger sample.

We uncovered critical misunderstandings in beliefs that could undermine the effectiveness of the epa's messages. For example, a sizable proportion of the public believes that radon contamination is permanent and does not go away. This misconception presumably results from an inappropriate inference based on knowledge about chemical contaminants or long-lived radioisotopes. The first version of the epa's "Citizen's Guide to Radon" did not discuss this issue. Based in part on our findings, the latest version addresses it explicitly.

The objective of risk communication is to provide people with a basis for making an informed decision; any effective message must contain information that helps them in that task. With former doctoral students Ann Bostrom, now at the Georgia Institute of Technology, and Cynthia J. Atman, now at the University of Pittsburgh, we used our method to develop two brochures about radon and compared their effectiveness with that of the epa's first version. When we asked people to recall simple facts, they did equally well with all three brochures. But when faced with tasks that required inference—advising a neighbor with a high radon reading on what to do—people who received our literature dramatically outperformed those who received the epa material.

We have found similar misperceptions in other areas, say, climatic change. Only a relatively small proportion of people associate energy use and carbon dioxide emissions with global warming. Many believe the hole in the ozone layer is the factor most likely to lead to global warming, although in fact the two issues are only loosely connected. Some also think launches of spacecraft are the major contributor to holes in the ozone layer. (Willett Kempton of the University of Delaware has found very similar perceptions.) The essence of good risk communication is very simple: learn what people already believe, tailor the communication to this knowledge and to the decisions people face and then subject the resulting message to careful empirical evaluation. Yet almost no one communicates risks to the public in this fashion. People get their information in fragmentary bits through a press that often does not understand technical details and often chooses to emphasize the sensational. Those trying to convey information are generally either advocates promoting a particular agenda or regulators who sometimes fail either to do their homework or to take a sufficiently broad perspective on the risks they manage. The surprise is not that opinion on hazards may undergo wide swings or may sometimes force silly or inefficient outcomes. It is that the public does as well as it does.

Indeed, when people are given balanced information and enough time to reflect on it, they can do a remarkably good job of deciding what problems are important and of systematically addressing decisions about risks. I conducted studies with Gordon Hester (then a doctoral student, now at the Electric Power Research Institute) in which we asked opinion leaders—a teacher, a state highway patrolman, a bank manager and so on—to play the role of a citizens’ board advising the governor of Pennsylvania on the siting of high-voltage electric transmission lines. We asked the groups to focus particularly on the controversial problem of health risks from electric and magnetic fields emanating from transmission lines. We gave them detailed background information and a list of specific questions. Working mostly on their own, over a period of about a day and a half (with pay), the groups structured policy problems and prepared advice in a fashion that would be a credit to many consulting firms.

If anyone should be faulted for the poor quality of responses to risk, it is probably not the public but rather risk managers in government and industry. First, regulators have generally adopted a short-term perspective focused on taking action quickly rather than investing in the research needed to improve understanding of particular hazards in the future. This focus is especially evident in regulations that have been formulated to ensure the safety of the environment, workplace and consumer products.

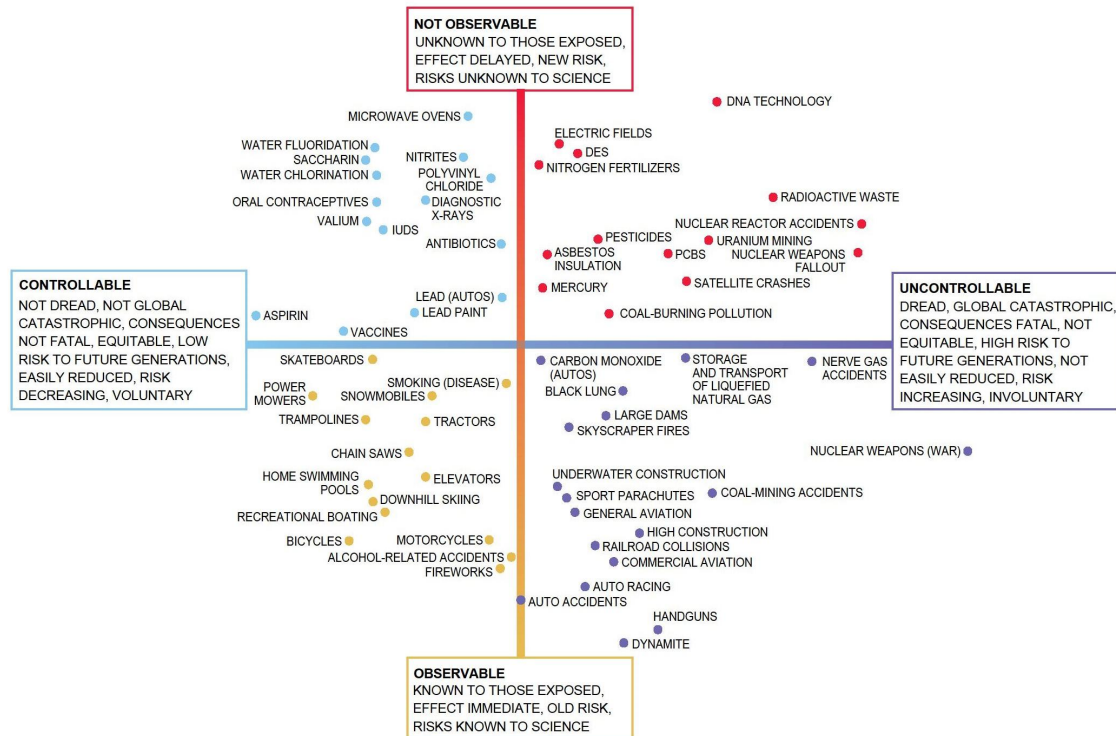
Second, these officials have often adopted too narrow an outlook on the risks they manage. Sometimes attempts to reduce one risk (burns from flammable children’s pajamas) have created others (the increased chance of cancer from fireproofing chemicals).

In some instances, regulators have ignored large risks while attacking smaller ones with vigor. Biologist Bruce Ames of Berkeley has argued persuasively that government risk managers have invested enormous resources in controlling selected artificial carcinogens while ignoring natural ones that may contribute far more to the total risk for human cancer.

Third, government risk managers do not generally set up institutions for learning from experience. Too often adversarial procedures mix attempts to figure out what has happened in an incident with the assignment of blame. As a result, valuable safety-related insights may either be missed or sealed away from the public eye. Civilian aviation, in contrast, has benefited extensively from accident investigations by the National Transportation Safety Board. The board does its work in isolation from arguments about liability; its results are widely published and have contributed measurably to improving air safety.

Many regulators are probably also too quick to look for single global solutions to risk problems. Experimenting with multiple solutions to see which ones work best is a strategy that deserves far more attention than it has received. With 50 states in a federal system, the U.S. has a natural opportunity to run such experiments.

Finally, risk managers have not been sufficiently inventive in developing arrangements that permit citizens to become involved in decision making in a significant and constructive way, working with experts and with adequate time and access to information. Although there are provisions for public hearings in the licensing process for

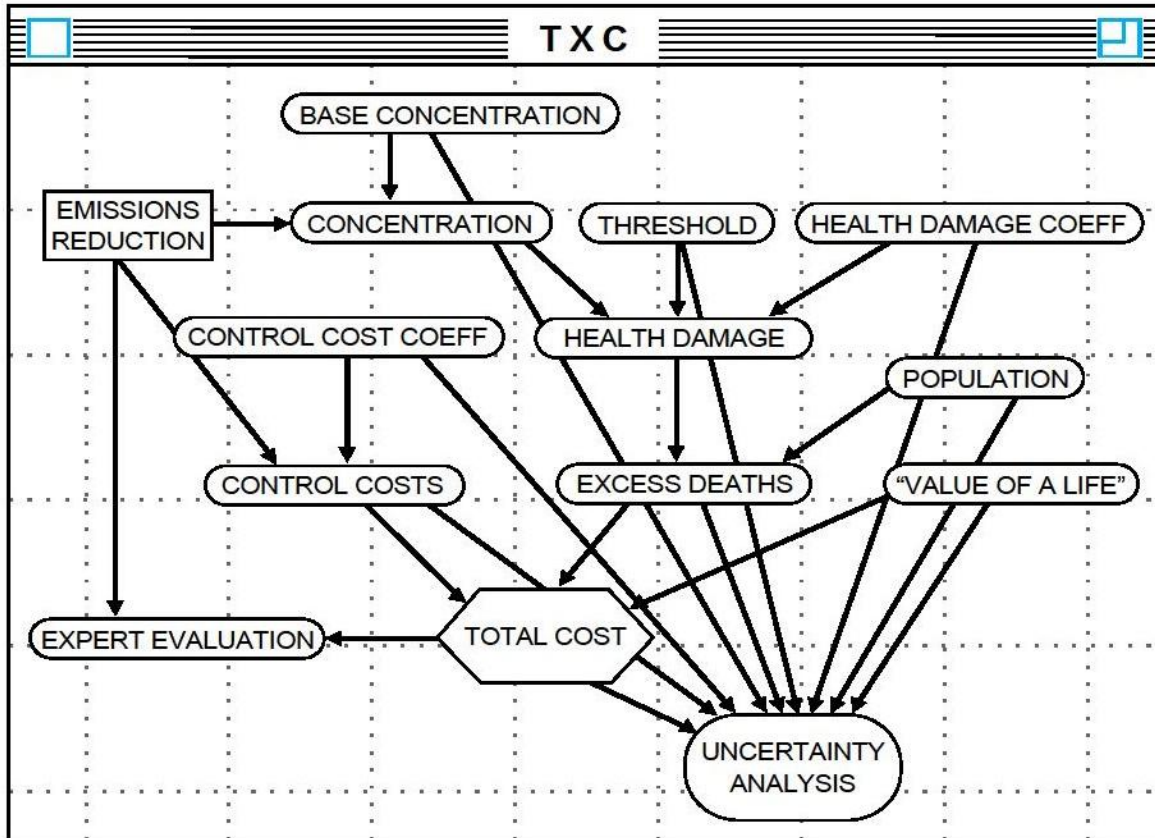


RISK SPACE has axes that correspond roughly to a hazard’s “dreadfulness” and to the degree to which it is understood. Risks in the upper right quadrant of this space are most likely to provoke calls for government regulation.

nuclear reactors or the siting of hazardous waste repositories, the process rarely allows for reasoned discussion, and input usually comes too late to have any effect on the set of alternatives under consideration.

Thomas Jefferson was right: the best strategy for assuring the general welfare in a democracy is a well-informed electorate. If the U.S. and other nations want better, more reasoned social decisions about risk, they need to take steps to enhance public understanding. They must also provide institutions whereby citizens and their representatives can devote attention to risk management decisions. This will not preclude the occasional absurd outcome, but neither does any other way of making decisions. Moreover, appropriate public involvement should go a long way toward eliminating the confrontational tone that has become so common in the risk management process.

Risk Analysis in Action



BLOCKS in the diagram above can be expanded to call up a window containing graphs and tables for their assumptions, equations and probability distributions.

Uncertainty is a central element of most problems involving risk. Analysts today have a number of software tools that incorporate the effects of uncertainty. These tools can show the logical consequences of a particular set of risk assumptions and rules for making decisions about it. One such system is Demos, developed by Max Henrion of Lumina Decision Systems in Palo Alto, Calif.

To see how the process works, consider a hypothetical chemical pollutant, “TXC.” To simplify matters, assume that the entire population at risk (30 million people) is exposed to the same dose—this makes a model of exposure processes unnecessary. The next step is to construct a function that describes the risk associated with any given exposure level—for example, a linear dose-response function, possibly with a threshold below which there is no danger.

Given this information, Demos can estimate the number of excess deaths caused every year by TXC exposure. According to the resulting cumulative probability distribution, there is about a 30 percent chance that no one dies, about a 50 percent chance that fewer than 100 people die each year and about a 10 percent chance that more than 1,000 die.

Meanwhile, for a price, pollution controls can reduce the concentration of TXC. (The cost of achieving any given reduction, like the danger of exposure, is determined by consultation with experts.) To choose a level of pollution control that minimizes total social costs, one must first decide how much society is willing to invest to prevent mortality. The upper and lower bounds in this example are \$300,000 and \$3 million per death averted. (Picking such numbers is a value judgment; in practice, a crucial part of the analysis would be to find out how sensitive the results are to the dollar values placed on life or health.)

Net social costs, in this model, are simply the sum of control costs and mortality. At \$300,000 per death averted, their most likely value reaches a minimum when TXC emissions are reduced by 55 percent. At \$3 million, the optimum reduction is about 88 percent.

Demos can also calculate a form of correlation between each of the input variables and total costs. Strong correlations indicate variables that contribute significantly to the uncertainty in the final cost estimate. At low levels of pollution control, possible variations in the slope of the damage function, in the location of the threshold and in the base concentration of the pollutant contribute the most to total uncertainty. At very high levels of control, in contrast, almost all the uncertainty derives from unknowns in the cost of controlling emissions.

Finally, Demos can compute the difference in expected cost between the optimal decision based on current information and that given perfect information—that is, the benefit of removing all uncertainties from the calculations. This is known in decision analysis as the expected value of perfect information; it is an upper bound on the value of research. If averting a single death is worth \$300,000 to society, this value is \$38 million a year; if averting a death is worth \$3 million, it is \$71 million a year.

Although tools such as Demos put quantitative risk analysis within reach of any group with a personal computer, using them properly requires substantial education. My colleagues and I found that a group of first-year engineering doctoral students first exposed to Demos tended to ignore possible correlations among variables, thus seriously overestimating the uncertainty of their results.

Further Reading

Rational Choice in an Uncertain World. Robyn M. Dawes. Harcourt Brace Jovanovich, 1988.

Readings in Risk. Edited by Theodore S. Glickman and Michael Gough. Resources for the Future, 1990.

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

Communicating Risk to the Public. M. Granger Morgan, Baruch Fischhoff, Ann Bostrom, Lester Lave and Cynthia J. Atman in *Environmental Science and Technology*, Vol. 26, No. 11, pages 2048-2056; November 1992.

Risk Analysis. Publication of the Society for Risk Analysis, published quarterly by Plenum Publishing.

The Ted K Archive

M. Granger Morgan

Risk Analysis and Management

Inadequate approaches to handling risks may result in bad policy. Fortunately,
rational techniques for assessment now exist

July 1993

Scientific American, July 1993.

www.thetedkarchive.com