Cambridge University Press
978-0-521-36542-0 - Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis
M. Granger Morgan and Max Henrion
Excerpt
More information

# 1 Introduction

> To know one's ignorance is the best part of knowledge.
> Lao Tzu, *The Tao*, no. 71

Life is full of uncertainties. Most of us have learned to live comfortably with day-to-day uncertainties and to make choices and decisions in their presence. We have evolved cognitive heuristics and developed strategies, technologies, and institutions such as weather reports, pocket-sized raincoats, and insurance to accommodate or compensate for the effects of uncertainty. Looked at with care, these heuristics and strategies do not always perform as well as we would like (Dawes, 1988). When our cognitive processes for dealing with uncertainty introduce error or bias into our judgments we are often unable to detect the fact. When things go seriously wrong we may not be around to learn the lesson – or we may still be unable to detect that the problem came from faulty processing of uncertain information. Thus, we muddle through – often doing quite well, occasionally getting into serious trouble.

Of course, uncertainty is not limited to our private lives. It also occurs in larger and more public situations. Frequently in public discussion, policy analysis, regulatory decision making and other contexts, we proceed as if we understand and can predict the world precisely. While a moment's reflection is sufficient to persuade anyone that this is not true, a number of political, behavioral, and analytical factors combine to promote the continuation of this practice.

This book is about dealing with scientific and technical uncertainty in risk analysis and in other forms of quantitative policy analysis and policy-focused research. Until recently this uncertainty has been treated in much the same way we have dealt with other uncertainties in our private and public lives. However, the past decade has seen a growing recognition that policies that ignore uncertainty about technology, and about the physical world, often lead in the long run to unsatisfactory technical, social, and political outcomes. Recent growth in interest, understanding, and technical skill in the field of risk analysis and assessment has worked to promote this change. By definition, risk involves an "exposure to a *chance* of injury or loss" (Random House, 1966). The fact that risk inherently involves chance or probability leads directly to a need to describe and deal with uncertainty.

The result has been a rapid rise in interest in techniques whose practical application was originally pioneered by workers in decision analysis like Howard Raiffa (Raiffa and Schlaifer, 1961; Raiffa, 1968) at Harvard and Ronald

Howard (1966) at Stanford. Good early examples of applications include the work of Jack Grayson (1960) on oil wildcatting; work by Ron Howard, Jim Matheson, and Warner North on controlling the risks posed by hurricanes (Howard, Matheson, and North, 1972) and work by Warner North and co-authors on such diverse problems as spacecraft biocontamination of Mars (North, Judd, and Pezier, 1973) and the risks of wildfires in the Santa Monica Mountains (North, Offensend, and Smart, 1975). Over the past decade, many new investigators have begun to work on theoretical and applied problems of uncertainty. A number of private sector decision makers, including several large corporations, now routinely employ decision analytic techniques that incorporate a treatment of uncertainty in their corporate planning and decision making. Several federal agencies including the U.S. Nuclear Regulatory Commission, the U.S. Department of Energy, and the U.S. Environmental Protection Agency have begun to address the problems of incorporating an explicit treatment of scientific and technical uncertainty in their analysis and regulatory decision making. Chapter 2 provides three examples of recent milestones in such public sector applications.

### 1.1. Does Uncertainty Really Matter?

The examples in the next chapter make it clear that dealing with uncertainty in policy analysis and policy-focused research has, at the very least, become "fashionable." On the other hand, for thousands of years our own and other societies have been muddling along, making decisions with less than complete attention to the associated uncertainties. Hence it seems only reasonable to ask whether, despite the current interest, uncertainty actually matters very much. In Chapter 3 and again in greater technical detail in Chapter 12, we explore this question at some length and conclude – not surprisingly, given this book's title – that, for quite a variety of reasons, uncertainty does matter and should not be ignored.

The detailed arguments can wait for later. Why, in a nutshell – in language that a casually interested airplane seatmate would understand – does technical uncertainty in risk assessment and other forms of policy research and analysis really matter? We could try an argument by analogy. Natural scientists are expected as a matter of course to include an estimate of the probable error when they report the value of quantities they have measured. The uncertainties involved in most quantitative policy analyses are much greater than those involved in work in the natural sciences. So policy analysts should report their uncertainties too.

Such an argument would probably satisfy some seatmates. On the other hand, the perceptive seatmate might insist on more specific arguments. In that case, we would offer three:

Cambridge University Press
978-0-521-36542-0 - Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis
M. Granger Morgan and Max Henrion
Excerpt
More information

1. A central purpose of policy research and policy analysis is to help identify the important factors and the sources of disagreement in a problem, and to help anticipate the unexpected. An explicit treatment of uncertainty forces us think more carefully about such matters, helps us identify which factors are most and least important, and helps us plan for contingencies or hedge our bets.
2. Increasingly we must rely on experts when we make decisions. It is often hard to be sure we understand exactly what they are telling us. It is harder still to know what to do when different experts appear to be telling us different things. If we insist they tell us about the uncertainty of their judgments, we will be clearer about how much they think they know and whether they really disagree.
3. Rarely is any problem solved once and for all. Problems have a way of resurfacing. The details may change but the basic problems keep coming back again and again. Sometimes we would like to be able to use, or adapt, policy analyses that have been done in the past to help with the problems of the moment. This is much easier to do when the uncertainties of the past work have been carefully described, because then we can have greater confidence that we are using the earlier work in an appropriate way.

The technical details of characterizing and dealing with uncertainty in policy research and analysis can get fairly complicated. It is important to remember that these *are* details and not lose sight of the big picture. Both people who commission and people who perform analysis that deals with uncertainty should from time to time return to the three broad arguments just listed and ask, "Is this really what our analysis is doing?" When the answer is not clearly yes, the time has come for some careful rethinking.

## 1.2. Outline of the Book

We close this introductory chapter with a brief outline of the balance of the book. For readers unfamiliar with the topic, Chapter 2 motivates the chapters that follow by presenting three recent and visible examples of risk and policy analyses in which uncertainty played an important role. Chapters 3 through 10 cover the central issues of the book. The final two chapters provide more specialized details. Chapter 3 presents a general overview of the subject of quantitative policy research and analysis and explores where and how the issue of uncertainty fits in this broader context. Chapter 4, "The Nature and Sources of Uncertainty," discusses various sources of uncertainty in some detail and then develops a general taxonomy of the kinds of quantities that can enter in policy analysis and the way in which to treat uncertainty about these quantities. There is a standard set of techniques for dealing with uncertainty in situations in which large amounts of relevant historical data, such as test results and time series, are available. Chapter 5, "Probability Distributions and Statistical Estimation," briefly surveys many of these techniques. This chapter, written principally by our colleague Mitchell Small, should serve to refresh the memory of those who have studied such techniques in the past. Mathematically knowledgeable readers who have had little or no training in probability and statistics should

Cambridge University Press
978-0-521-36542-0 - Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis
M. Granger Morgan and Max Henrion
Excerpt
More information

find this chapter a useful overview of key portions of the field. However, before taking action on the basis of this overview, readers without previous training in probability and statistics are urged to reflect seriously on the adage that "a little knowledge can be a dangerous thing."

Although it is important to make use of whatever good evidence is available, for many problems addressed by policy analysis the evidence available is insufficient to allow for the application of classical statistical techniques of the kind reviewed in Chapter 5. In these circumstances, about all one can do is resort to the use of "expert judgment," as explored in Chapters 6 and 7. The past fifteen years have witnessed very considerable progress in our understanding of the way in which people think about and make judgments in the presence of uncertainty. This literature is reviewed in Chapter 6, and Chapter 7 discusses the actual mechanics of eliciting subjective probabilistic judgments from experts.

Chapter 8, "The Propagation and Analysis of Uncertainty," is a fairly complete discussion of the various techniques for propagating uncertainty through quantitative policy models and for analyzing the implications of uncertainty in such models.

Representing, modeling, and analyzing uncertainty are of limited use if one cannot effectively communicate to other people the results of these efforts. Probably the most efficient means for such communication is through pictures. In Chapter 9, "The Graphic Communication of Uncertainty," we explore this subject both from the perspective of communication to technical people and from the perspective of communication to semitechnical and nontechnical people.

Most techniques discussed in Chapter 8 require the use of a computer. Not all computers are equally helpful in quantitative policy analysis. Conventional computer environments that support standard procedural languages such as FORTRAN often do not make it possible to engage in many of the activities identified in Chapter 3 as being important in "good" policy analysis. Chapter 10 describes an experimental computer environment that we have constructed, discusses the various evaluative experiments we have run, and outlines some of the more general insights we have gained while performing this research.

The final two chapters are more specialized. Chapter 11 explores a number of problems related to very large models. Chapter 12 is a decision theoretic discussion of the value of including uncertainty in an analysis.

## References

Dawes, R. W. (1988). *Rational Choice in an Uncertain World*, Harcourt Brace & Jovanovich, New York.

Grayson, C. J. (1960). *Decisions Under Uncertainty: Drilling Decisions by Oil and Gas Operators*, Harvard Business School, Plimpton Press, Cambridge, Mass.

Howard, R. A. (1966). "Decision Analysis: Applied Decision Theory," *Proceedings of the Fourth International Conference on Operational Research*, ed. D. B. Hertz and J. Melese, pp. 55–71, Wiley-Interscience, New York.

Cambridge University Press
978-0-521-36542-0 - Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis
M. Granger Morgan and Max Henrion
Excerpt
More information

Howard, R. A., Matheson, J. E., and North, D. W. (1972). The Decision to Seed Hurricanes, *Science*, 176, (February 23):1191–1202.

North, D. W., Judd, B. R., and Pezier, J. P. (1973). "New Methodology for Assessing the Probability of Contaminating Mars," paper presented to the XVII meeting of COSPAR, São Paulo, Brazil, (COSPAR Identification V.4.3).

North, D. W., Offensend, F. L. and Smart, C. N. (1975). "Planning Wildfire Protection for the Santa Monica Mountains," *Fire Journal* (January).

Raiffa, H. (1968). *Decision Analysis: Introductory Lectures on Choices Under Uncertainty*, Addison-Wesley, Reading, Mass.

Raiffa, H., and Schlaifer, R. O. (1961). *Applied Statistical Decision Theory*, Harvard Business School, Cambridge, Mass.

Random House (1966). *The Random House Dictionary of the English Language*, ed. Jess Stein, Random House, New York.

# 2 Recent Milestones

> Probabilities direct the conduct of the wise man.
> Cicero, *De Natura Deorum*,
> Book 1, chap. 5, sec. 12

Although a considerable theoretical literature and a number of small scale applications of techniques for dealing with uncertainty in policy analysis and policy focused research have been around for a couple of decades, larger applications to major policy problems are a more recent phenomenon. Because many of the ideas and techniques involved are new, even to members of the technical community, their introduction into policy circles has been uneven and accompanied by a variety of mistakes and false starts. Nevertheless, there have now been a number of important applications, and several U.S. federal agencies have become seriously committed to the continued development and use of these techniques. For readers unfamiliar with these developments we briefly motivate the discussions that follow with three case examples that involve techniques for incorporating an explicit treatment of uncertainty in (1) estimates of the safety of light-water nuclear reactors; (2) the regulatory analysis of common ("criteria") air pollutants; and (3) estimates of the probable impacts on the ozone layer of continued release of chlorofluorocarbons.

## 2.1. Reactor Safety

One of the earliest large-scale studies to employ a formal treatment of uncertainty was the Reactor Safety Study, NUREG-75/014 (WASH-1400), generally known as the "Rasmussen report" (Rasmussen et al., 1975). In the early 1970s what was then the U.S. Atomic Energy Commission (AEC) asked Norman C. Rasmussen, a professor of nuclear engineering at the Massachusetts Institute of Technology, to undertake a quantitative study of the safety of light-water reactors. In order to make the task concrete, two specific commercial light-water power reactors, one a pressurized water reactor and one a boiling water reactor, were selected for study. Rasmussen assembled a team of roughly sixty people, who undertook to identify and formally describe, in terms of event trees, the various scenarios they believed might lead to major accidents in each of the two reactors studied. Fault trees were developed to estimate the probabilities of the various events. A combination of historical evidence from the nuclear and other industries, together with expert judgment, were used to construct the probability estimates, most of which were taken to be log-normally distributed. The probabilities of the various accident scenarios were then estimated through stochastic simulation. To assess the consequence of these accidents, models of exposure process for released radioactive contaminants were employed. Models of effects processes were added to assess the immediate and delayed health consequences to populations exposed to the releases. These exposure and effects models did not incorporate as much attention to the treatment of uncertainty as did the reactor accident models. Finally, an attempt was made to compare the resulting risks with a variety of natural and other man-made hazards. In summary:

> The Rasmussen group estimated the overall probability of a sequence of events leading to a core melt as one chance in 20,000 reactor-years of operation; they found the average core melt would itself not present a major threat to the public health and safety, and they concluded (in a widely quoted and much criticized statement) that the likelihood of an average citizen's being killed in a reactor accident is about the same as the chance of his being killed by a falling meteorite. (Lewis, 1980)

The Rasmussen report was greeted with both great interest and substantial criticism. Some of the criticisms involved valid technical concerns, some were adversarial reactions motivated by opposition to nuclear power. To obtain an independent evaluation and deal with the criticisms the U.S. Nuclear Regulatory Commission (NRC) appointed a second committee, the Risk Assessment Review Group under the chairmanship of another academic, Harold W. Lewis, a professor of physics at the University of California at Santa Barbara. The resulting "Lewis report" (Lewis et al., 1975) confirmed many of the technical criticisms of the Rasmussen report, including criticisms about the treatment of multiple failures resulting from a common cause, and the ways in which uncertainties were represented, propagated, and interpreted. It drew attention to the difficulties of incorporating various human elements such as human adaptability during accidents; to the pervasive regulatory influence in the choice of uncertain parameters; to the "inscrutable" nature of the written report; and to the fact that the executive summary did not provide a good description of the contents of the report. However, despite such concerns about the details, the report concluded that the techniques developed and demonstrated in the Rasmussen study were, in Lewis's words,

> extremely valuable and should be far more widely applied in the process of regulating the nuclear industry. Such probabilistic techniques, which provide guidance on the important issues in reactor safety, would be helpful in determining the priorities of the NRC both in its safety-research program and in the deployment of its regulatory and inspection resources. (Lewis, 1980)

At about the same time, the American Physical Study released a study of light-water reactor safety that they had conducted under National Science Foundation (NSF) and AEC support[1] (APS, 1975). Although this latter study was not undertaken as a specific evaluation of the Rasmussen report, the APS group did

---

[1] There was considerable overlap in membership of this study group and the Risk Assessment Review Group. Lewis was chairman of both groups.

review the Rasmussen report in considerable detail. On the validity of such studies they observed:

> It is difficult to quantify accurately the probabilities that any accident-initiating event might occur. Many aspects need to be better understood through experience and research before such calculations are tractable . . . we recognize that the event-tree and fault-tree approach can have merit in highlighting *relative* strengths and weaknesses of reactor systems, particularly through comparisons of different sequences of reactor behavior. However, based on our experience with problems of this nature involving very low probabilities, we do not now have confidence in the presently calculated absolute values of the probabilities of the various branches. (APS, 1975)

Despite these reservations, the conclusions of the APS group call for significantly expanded research on probabilistic techniques including checks against empirically observed rates for small accidents and "parametric studies of phenomena which are ill understood in the identified sequences" and the adoption of quantitatively stated safety goals for reactor performance. (APS, 1975)

Although it had various shortcomings, the Rasmussen report got many things right. For example, it concluded that transients, small loss-of-coolant accidents, and human errors are important contributors to the overall risks of nuclear power. Lewis points out that "these three items . . . were the central features of the Three Mile Island accident" (Lewis, 1980). While they have continued to undergo refinement, the general techniques pioneered in the Rasmussen report are now used widely and routinely in a variety of applications by the NRC, by the atomic energy agencies of many other nations, and by the nuclear industry, including applications such as the design of new reactors (Gonzalez et al., 1987).

The Rasmussen study, and the several dozen probabilistic reactor safety studies that have followed it, played a substantial role in stimulating the Nuclear Regulatory Commission to develop probabilistic "safety objectives" in addition to their traditional "qualitative safety goals." After a series of workshops, staff studies, and public comments, in August 1986 the NRC promulgated the following specific objectives:

- The risks to an average individual in the vicinity of a nuclear power plant of prompt fatalities that might result from reactor accidents should not exceed one-tenth of one percent (0.1 percent) of the sum of prompt fatalities resulting from other accidents to which members of the U.S. population are generally exposed.
- The risk to the population in the area near a nuclear power plant of cancer fatalities that might result from nuclear power plant operation should not exceed one-tenth of one percent (0.1 percent) of the sum of cancer fatality risks resulting from all other causes (NRC, 1986).

Although they did not promulgate it, they also indicated their intention to consider a quantitative objective for accident rates:

> Consistent with the traditional defense-in-depth approach and the accident mitigation philosophy requiring reliable performance of containment systems, the overall mean frequency of a large release of radioactive materials to the environment from a reactor accident should be less than 1 in 1,000,000 per year of reactor operation (NRC, 1986).

The commission recognized the probabilistic nature of any such estimates. In the context of the safety objectives they explained that:

> The Commission is aware that uncertainties are not caused by the use of quantitative methodology in decision making but are merely highlighted through use of the quantification process. Confidence in the use of probabilistic and risk assessment techniques has steadily improved … . In fact, through the quantitative techniques, important uncertainties have been and continue to be brought into better focus and may even be reduced compared with those that would remain with a sole reliance on deterministic decision making. To the extent practicable the Commission intends to ensure that the quantitative techniques used for regulatory decision making take into account the potential uncertainties that exist so that an estimate can be made on the confidence level to be ascribed to the quantitative results (NRC, 1986).

The commission goes on to explain that for the purposes of implementing their quantitative safety objectives, they intend to use mean values, but that they intend to remain alert to the ranges associated with the estimates and with the phenomenology that produces these ranges.

The process of implementing these objectives, of developing benefit–cost techniques to evaluate proposed changes, and of adopting the proposed accident rate guideline has proceeded slowly and has stimulated all kinds of activity and controversy (Cave, 1987). The proposal has been made to require all U.S. nuclear plants to provide a probabilistic risk assessment that shows their expected performance with respect to the quantitative safety objectives. Because the commission developed the objectives with existing assessments in mind, it is perhaps not surprising that assessments completed to date of a number of U.S. plants show the objectives being met (Cave, 1987). Of course, most of these assessments contain significant model assumptions, some of whose consequences have not yet been sufficiently explored. As alternative model assumptions are made in future assessments, compliance may not always be so automatic.

Although these and other activities related to the design and regulation of nuclear power have substantially stimulated the development and wider use of techniques for dealing with uncertainty in quantitative policy studies, they have certainly not resolved the basic dilemmas that the nuclear industry faces. That, of course, is because these dilemmas involve issues far more fundamental than accurately describing uncertain accident rates and consequences.

## 2.2. Air Pollution

The Clean Air Act requires the U.S. Environmental Protection Agency to establish, and periodically revise, national ambient air quality standards for a number of ubiquitous or “criteria” air pollutants. Until the late 1970s the EPA approached this task by first reviewing the relevant atmospheric and biological literature and then applying what was basically a seat-of-the-pants approach that relied on agency personnel to make judgments regarding complex health issues.

Cambridge University Press
978-0-521-36542-0 - Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis
M. Granger Morgan and Max Henrion
Excerpt
More information

Later this approach was expanded to include expert advice from outside scientific review panels. This process mixed scientific, economic, and value judgments without differentiating them, and many of the participants were unaware they were being made. Understandably, the EPA's Office of Air Quality Planning and Standards (OAQPS) was uncomfortable with these procedures. Thus, in 1977 when OAQPS began the process of reevaluating the ambient standard for ozone, they opted to pursue a dual strategy. In addition to their traditional approach the office put Thomas Feagans and William Biller to work on developing a formal decision analytic approach that would incorporate the very considerable scientific uncertainties through the use of elicited expert subjective judgments (Feagans and Biller, 1978, 1981).

The Feagans–Biller work suffered from a number of serious problems in both its conception and its execution. The process of promulgating any new national ambient air quality standard is contentious, and the ozone standard proved particularly controversial. In the course of this controversy the Feagans–Biller work was severely criticized by both outside groups, such as the American Petroleum Institute (API, 1978), and by several EPA advisory panels. Although it had initially tried to use the Feagans–Biller work as part of the process of setting the new ozone standard, in the face of this criticism the EPA backed off and subsequently argued that the Feagans–Biller work had been only experimental in nature. A more detailed, if not entirely balanced, history of this rocky beginning can be found in Marraro (1982). In the aftermath of this fight, a special subcommittee of the EPA Science Advisory Board, the Committee on Health Risk Assessment, was established to assist OAQPS in developing a satisfactory approach.[2] Following advice from this committee the agency gradually abandoned its earlier formulation of the problem and turned to the community of experienced decision analysts for ideas and advice. In April 1980, OAQPS convened a "brainstorming" session at which a number of these experts presented concept papers they had developed under a series of small contracts.[3] At about the same time, a contract was let with Thomas Wallsten of the University of North Carolina to undertake a review of the psychological literature on expert elicitation (Wallsten and Bendescu, 1980, 1983). After reviewing the results of the brainstorming session, larger follow-up contracts were let to two groups – the first under the general leadership of Miley Merkhofer at SRI International (Smith, McNamee, and Merkhofer, 1982) and the second under the joint leadership of Rakesh Sarin of UCLA,

---

[2] One of us, Granger Morgan, served on this committee and the several subsequent EPA/SAB committees that have overseen OAQPS's work in this area.

[3] Authors who contributed to this process included Howard Raiffa and Richard Zeckhauser; Richard de Neufville and Marie-Elizabeth Pate (Cornell); Chris Whipple and Baruch Fischhoff; Kenneth Manton, H. O. Hartley, and Max Woodbury; Robert Winkler and Rakesh Sarin; and Miley Merkhofer. For a summary, see EPA (1981).