

Chapter 1

Introduction and Overview

G.H. Hardy: Here lies Hardy, with no apologies.

1.1 PREAMBLE: WHAT DO ‘RELIABILITY’, ‘RISK’ AND ‘ROBUSTNESS’ MEAN?

Words such as ‘credibility’, ‘hazard’, ‘integrity’, ‘reliability’, ‘risk’, ‘robustness’ and ‘survivability’ have now become very much a part of our daily vocabulary. For instance, the derogatory term **unreliable** is used to describe the undependable behavior of an individual or an item, whereas the cautionary term **risky** is used to warn of possible exposure to an adverse consequence. The term **survival** is generally used in biomedical contexts and is intended to convey the possibility of overcoming a life-threatening situation or a disease. **Robustness** encapsulates the feature of the persistence of some attribute in the presence of an insult, such as a shock, or an unexpectedly large change, such as a surge in electrical power, or an encounter with an unexpectedly large (or small) observation. Thus robustness imparts the attribute of reliability to a physical or a biological unit, and sometimes even to a mathematical or a statistical procedure. In what follows, we point out that all the above terms convey notions that are intertwined, and thus, in principle, they tend to be used interchangeably. Our choice of the words ‘reliability’ and ‘risk’ in the title of this book reflects their common usage.

It is often the case, even among engineers and scientists, that the above terminology is purely conversational, and is intended to convey an intuitive feel. When such is the case, there is little need to be specific. Often, however, with our increasing reliance on technology, and for decisions pertaining to the use of a technology, we are required to be precise. This has resulted in efforts to sharpen the notions of risk and reliability and to quantify them. Quantification is required for normative decision making, especially decisions pertaining to our safety and well-being; some examples are mentioned in section 1.4. When quantified measures of risk are coupled with normative decision making, it is called **risk management** (cf. The National Research Council’s Report on “Improving the Continued Airworthiness of Civil Aircraft”, 1998).

Historically, the need for quantifying risk pre-dates normative decision making. It goes back to the days of Huygens (1629–1695), who was motivated by problems of annuities. In the early 1930s, it was problems in commerce and insurance that sustained an interest in this topic. During

2 INTRODUCTION AND OVERVIEW

the 1960s and the 1970s, quantified measures of reliability were needed to satisfy specifications for government acquisitions, mostly in aerospace and defense; and quantified measures of risk were needed for regulation, mostly drug approval, and for matters of public policy, such as reactor safety. During the 1980s, pressures of consumerism, competitiveness and litigation have forced manufacturers and service organizations to use quantified measures of reliability for specifying assurances and designing warranties. The coming era appears to be one of ensuring infrastructure integrity, infrastructure protection and with the advent of test-ban treaties, the stewardship of nuclear weapons stockpiles. Here again, quantified measures of reliability are poised to play a signal role.

The above developments have given birth to the general topic of **risk analysis**, which in the context of engineering applications takes the form of **reliability analysis**, and in the context of biomedicine, **survival analysis**. These two scenarios have provided most of the applications and case histories. The 1990s have also witnessed the use of risk management in business and finance – for acquisitions, bond pricing, mergers, and the trading of options – and in political science for matters of disarmament and national security. The applications mentioned above, be they for the design of earthquake-resistant structures, or for the approval of medical procedures, have one feature in common: they all pertain to situations of **uncertainty**, and it is this common theme of uncertainty that paves the way for their unified treatment. Uncertainty is about the occurrence of an undesirable event, such as the failure of an item, or an adverse reaction to a drug, or some other loss. Since a conversational use of the words ‘reliability’ and ‘risk’ conveys an expression of uncertainty, it is the quantification of uncertainty that is, *de facto*, the quantification of reliability and risk. To summarize, reliability and risk analysis pertains to quantified measures of uncertainty about certain adverse events.¹ However, since quantified measures of uncertainty are only an intermediate step in the process of normative decision making, one may take a broader view and claim that reliability and risk analysis is simply methods for decision making under uncertainty. This is the point of view taken in ‘Risk: Analysis, Perception, and Management. Report of a *Royal Society Group*’ (1992).

The quantification of uncertainty is an age-old problem dating back to the days of Gioralimo Kardano (1501–1575) (cf. Gnedenko, 1993), and decision making under uncertainty can trace its roots to von Neumann and Morgenstern’s (1944) theory of games and economic behavior, if not to Daniel Bernoulli (1700–1782). It was Bernoulli who, in proposing a solution to the famous ‘St. Petersburg paradox’, introduced the idea of a utility, i.e. the consequence of each possible outcome in a situation of uncertainty. Thus, putting aside the matter of a focus on certain types of events, what is new and different about reliability and risk analysis, and why do we need another book devoted to this topic?

The answer to the first part of the above question is disappointing. It is that from a foundational point of view there is nothing special about problems in reliability and risk analysis that the existing paradigms used to quantify uncertainty cannot handle. The fundamental territory has been introduced, developed and explored by individuals bearing illustrious names like Bernoulli, De Moivre, de Finetti, Fermat, Huygens, Laplace, Pascal, and Poisson. I attempt to answer the second part of the question in the following sections, but I am unsure of success. This is because my main reason for writing this book is to articulate a way of conceptualizing the problems of reliability and risk analysis, and to use this conceptualization to develop a unified approach to quantify them. I warn the reader, however, that my point of view may not be acceptable to all, though my hope is that once the fundamentals driving this point of view are appreciated and understood – which I hope to do here – my position will be more palatable.

¹ Not to be considered as being synonymous with Heisenberg’s ‘uncertainty principle’, which says that in the quantum mechanical framework the error (uncertainty) in the measurement of position multiplied by the error (uncertainty) in time measurement must exceed a certain constant called Planck’s constant. This principle was enunciated by Werner Heisenberg in the mid-1920s.

1.2 OBJECTIVES AND PROSPECTIVE READERSHIP

The aim of this book is twofold. The first is to discuss a mathematical framework in which our uncertainty about certain adverse effects can be quantified, so that the notions of hazard, risk, reliability and survival can be discussed in a unified manner. The second aim is to describe several reliability and risk analysis techniques that have been developed under the framework alluded above. My intention is to focus strongly on the matter of how to think about reliability and risk, rather than to focus on particular methodologies. Since the quantification of uncertainty has been the subject of much debate, it is essential that the key arguments of this debate be reviewed, so that the point of view I adopt is put in its proper context. Thus we start off with an overview of the philosophical issues about the quantification of uncertainty, and decision making in the presence of uncertainty. The overview material should be familiar to most graduate students in probability and statistics; however, those who have not had exposure to Bayesian thinking may find it useful. The overview is followed by a description of the key ideas and methodologies for assessing reliability and risk. The latter material constitutes the bulk of my effort and should appeal to those with applied interests. However, the importance of the foundational material needs to be underscored; it sets the tone for the ensuing developments and provides a common ground for addressing the various applications.

Because of the current widespread interest in reliability, risk and uncertainty, the book should be of appeal to academics, students and practitioners in the mathematical, economic, environmental, biological and the engineering sciences. It has been developed while keeping in mind these multiple communities. The material here could possibly also be of some interest to quantitative philosophers and mathematically oriented specialists in the areas of medicine, finance, law, national security and public policy. However, by and large, the bulk of its readership would come from graduate students in engineering, systems analysis, operations research, biostatistics and statistics. With the above diversity of clientele, it is important to draw the reader's attention to one matter. Specifically, throughout this book, the uncertain event that I focus upon is the failure of items in biological and engineering systems, rather than, say, the occurrence of a financial or a strategic loss. This admission would appear to suggest that the material here may not be of relevance to risk analysts in business, finance and other such areas. This need not be true, because the manner in which I propose to quantify uncertainty is not restricted to a particular class of problems. The choice of applications and examples is determined by my experience, which necessarily is limited. All the same, I admit to the difficulty of using my approach for addressing risk problems that are basically communal or political. Matters pertaining to single issue campaigns involving extreme positions are not in the scope of our discussion.

1.3 RELIABILITY, RISK AND SURVIVAL: STATE-OF-THE-ART

Even though the first truly empirical mortality table was constructed as early as 1693 by Edmund Halley, formal material on survival analysis appeared in actuarial journals – mostly Scandinavian. The term 'hazard rate' seems to have originated there. Since the late 1960s the literature on reliability, risk and survival analysis has experienced an explosive growth. It has appeared in diverse and scattered sources, ranging from journals in philosophy, mathematics (predominantly statistics), biomedicine, engineering, law, finance, environment and public policy. In the last 30 years or so, there have been annual conferences and symposia devoted to the various aspects of these topics. More recently, journals that pertain exclusively to these subjects have also begun to appear. These conferences, journals and symposia have been sponsored by different professional groups, with different interests and different emphases.

The activities of the various interest groups have been unconnected because each emphasizes a particular type of an application, or a particular point of view or a particular style of analysis.

4 INTRODUCTION AND OVERVIEW

For example, risk analysis done by nuclear engineers and physicists is significantly different from survival analysis done by biostatisticians. The difference is due not just to the nature of their applications, but more so to their attitudes toward quantifying uncertainty. Physicists, being generally trained as objective scientists, have come to realize and to accept subjectivity in the sciences (cf. Schrodinger, 1947); consequently, their analyses of technological risks have incorporated subjective elements. By contrast, biostatisticians, being subjected to public scrutiny, have tended to be cautious and very factual with their analyses. Similarly, the type of reliability analysis done by, say, an electronics engineer differs from that done by an applied probabilist or a statistician. The former tends to emphasize the physics of failure and tends to downplay the mathematics of uncertainty; the latter two tend to do the opposite. Certainly there is a need for both, the physics of failure and the mathematics of uncertainty. A mathematical paradigm that can formally incorporate the physics of failure into the quantification of uncertainty would help integrate the activities of the two groups and produce results that would be more realistic. A goal of this book is to present a point of view that facilitates the above interplay.


A different type of situation seems to have arisen with software engineers working on reliability problems. They often do a credible job analyzing the causes of software failure, but then quantify their uncertainties using a myriad of analytical techniques, many of them ad hoc. This has caused much concern about the state-of-the-art of software risk assessment (cf. Statistical Software Engineering, 1996). Difficulties of this type have also arisen in other scenarios. A possible explanation is that most risk analysts tend to be subject matter specialists who concentrate on the science of the application and tend to accept analytical methodologies without an evaluation of their limitations and theoretical underpinnings.

A credible risk analysis requires an integration of the activities of the various groups, and one step toward achieving this goal would be to identify and agree upon a common theme under which the notions of risk, reliability and survival can be discussed and quantified, independent of the application. Our purpose here is to advocate such a theme and I endeavor to do so in a manner understandable by subject matter specialists who are mathematically mature; as a result, all risk analysts will have a common basis from which to start and a common goal to aim for.

1.4 RISK MANAGEMENT: A MOTIVATION FOR RISK ANALYSIS

We have seen that risk management is decision making under uncertainty using quantified measures of the latter. The purpose of this section is to elaborate on the above with a view toward providing a motivation for conducting risk analysis. A few scenarios are cited that are suitable candidates for risk management and, in one case, the steps by which one can proceed are outlined.

Most studies in risk management come into play because of the possible dangers faced whenever new technologies in engineering or medicine are proposed. New technologies advance the way we live, but occasionally at a price. In some cases, the price is unacceptable. The objective in risk management is to investigate the trade-off between the conveniences and the consequences, both tangible and intangible. The end result is often a binary decision to introduce, or not, a proposed technology. As an illustration, consider the following medical scenario: cholesterol-lowering drugs decrease the chances of heart attacks, but are expensive to administer, cause physical discomfort to the patient and often have side effects such as damage to the kidneys. If in an individual case, the possibility of the side effects materializing is small, if the patient is willing to bear the cost of the drug and is able to withstand its discomfort, then there is more to be gained by prescribing the drug, than by not. But to arrive at such a decision we need to quantify at least six uncertainties, four connected with the administration of the drug, and the remaining two if the drug is not administered (Figure 1.1). Risk analysis is the



Key:

- | | |
|---|--|
| p_i and q_i : Probabilities; | U_i : Utilities; |
| A : Administer Drug; | (A) : Do not Administer Drug; |
| $H.A$: Patient Suffers a Heart Attack; | $(H.A)$: Patient Avoids a Heart Attack; |
| $S.E$: Patient Suffers Side Effects; | $(S.E)$: Patient Avoids Side Effects. |

Figure 1.1 Decision tree for the cholesterol lowering drug problem.

process of quantifying such uncertainties. We also need to evaluate the patient’s **utilities**, i.e. the consequences, usually expressed as costs, associated with each of the above six uncertainties. The utilities should include the cost of administering the drug, a cost figure associated with the discomfort of side effects, cost figures attached to suffering a heart attack, and the rewards of avoiding one. The assessment of utilities is a very crucial and, perhaps, the most difficult step in risk analysis. It is best performed by individuals knowledgeable in economics and the behavioral sciences. Determining an individual’s utilities usually involves asking them to express preferences among different options. Clearly, assigning cost figures to the consequences of having a heart attack and to the merits of avoiding one is not a standard exercise, but one that nevertheless has to be addressed. The matter of utilities is discussed in some detail in section 2.8. Quantifying uncertainties is often a detailed task which, in our example, would start with the patient’s medical history and would involve tracing the causes of a heart attack and the drug’s side effects. A useful device that graphically portrays the causes and the sequences of events that lead to an undesirable event is a **fault tree**, also known as an **event tree**; see Barlow, Fussell, and Singpurwalla (1975) for examples on how to construct fault trees. The more detailed an accounting of the causes, the more credible is the quantification of uncertainty. Clearly, such a task would require the active participation of subject matter specialists, such as physicians and biochemists. The actual quantification should be done by someone knowledgeable in the calculus of uncertainties, which we hope that this book will help quantitative subject matter specialists to become. Thus, risk analysis is often a multi-disciplinary process involving participation by economists, mathematicians, social scientists, engineers and other subject matter specialists.

Figure 1.1 is a decision tree which shows the various steps that are involved in dealing with the cholesterol drug problem. The rectangle at the leftmost end of the tree denotes a decision node, which shows the two possible actions that a decision maker, usually the

physician, can take: \mathcal{A} – to administer the drug; or (\mathcal{A}) – not to administer it, the parentheses surrounding \mathcal{A} denote its complement. The circles denote the random nodes, which show the possible (unpredictable) outcomes that can occur under each action taken by the decision maker. The notation $\mathcal{H}\mathcal{A}$ denotes the event that the patient suffers a heart attack, whereas $(\mathcal{H}\mathcal{A})$ denotes the event that the patient escapes it. Similarly, $\mathcal{S}\mathcal{E}$ denotes the event that the patient experiences the drug's side effects, whereas $(\mathcal{S}\mathcal{E})$ denotes the event that the patient does not experience side effects. Observe that a heart attack can occur whether or not the drug is administered, but that there can be no side effects when the drug is not administered. At the right-hand end of the tree, we indicate the patient's utilities that are encountered with each of the six terminal branches of the tree; these are denoted by the U_i 's. Important to Figure 1.1 are the numerical values that quantify the uncertainties associated with the events describing each of the six terminal branches; these have been denoted by the p_i 's and the q_i 's. A focus of this book is to describe procedures by which such numerical values can be assigned. Once this has been done, standard decision theory (cf. Lindley, 1985, pp. 139–159) prescribes a procedure by which the decision to administer the drug, or not, can be made. More details about this are given in section 2.9. The decision would depend on the assessed values of the patient's utilities and the numerical values of the assessed uncertainties. These of course would vary from patient to patient.

A similar type of analysis should be used for analyzing the risk of introducing the recently proposed 'fly-by-wire airplanes', in which the control of aircraft is under the direction of a computer. The advantage of such airplanes is less reliance on pilots who are prone to human error. Their main disadvantage is the possible presence of a fatal flaw in the software which directs the computer. What are the chances of having such a flaw and, should there be one, what are the chances of encountering it during flight? A numerical assessment of these chances, together with an assessment of utilities, would enable us to decide whether or not to commission the fly-by-wire airplanes. A less daunting example, also in connection with airplane travel, pertains to the current trend by aircraft manufacturers to equip large passenger jets with only two engines rather than the usual four, which is known as ETOPS (for extended twin engine operations). Whereas most people would prefer to fly in aircraft equipped with four engines, it is possible that having only two engines lowers the stresses on the rest of the airplane, making its overall reliability better than that of an aircraft with four engines. Decisions of this type should be supported by a formal exercise in risk management; see Appendix E of the National Research Council's 1998 Report mentioned earlier. Also see sections 10.5 and 10.6, wherein we describe decision making for allocating reliability in systems design and for system selection (procurement), respectively. In the context of reliability, decision trees also come into play when one considers life testing and the design of life testing experiments (sections 5.5 and 5.6).

Indeed, if one of the main purposes of doing reliability analysis is to facilitate a good engineering design, and because design involves a trade-off among alternatives, one may take the view that the purpose of a reliability study is to help make sound decisions in the face of uncertainty.

1.5 BOOKS ON RELIABILITY, RISK AND SURVIVAL ANALYSIS

In section 1.1, we raised the issue about the need for another book on reliability and survival analysis. I have given the reader some hints about my aims here but have said little about what is currently available and how it differs from what is planned here. In what follows, a broad-brushed perspective on published material on these topics is given.

The existing books and monographs on reliability, risk and survival analysis can be classified into the following three categories: 1) works that are heavily focused on the subject matter details of a particular application, say nuclear reactor safety; 2) works that develop models for

quantifying uncertainties and which focus on the detailed mathematical structure of such models; and 3) works that emphasize statistical issues pertaining to the quantification of uncertainty and the treatment of data. The first category does not warrant concern vis-à-vis duplication because the material there does not advocate an overall theme for addressing a general class of problems, nor does it articulate any particular paradigm for thinking about problems in reliability and risk analysis. The treatment is generally on a case-by-case basis, and its greatest appeal is to practitioners whose interests are non-mathematical. In the second category, of which the two books by Barlow and Proschan (1965, 1981) are landmarks, the emphasis is on material that may be labeled ‘academic’. The authors refer to their work as a ‘mathematical theory of reliability’, and correctly so. The main handicap of this second category is that the initial uncertainties are treated as being quantified (via probability), and the emphasis is on how these initial uncertainties propagate. That is, the initial probabilities (be they objective, subjective or logical) are assumed known. The attitude there is more in keeping with the Russian school of probability, wherein the source and nature of initial probabilities are not a matter of concern. Furthermore, the mathematical theory is not integrated into the broader framework of risk management. That is, its place in the context of decision making under uncertainty has not been sufficiently articulated. For us here, the real overlap – if any – is with respect to the third category, which also happens to be the biggest. Representatives of this group that have a focus toward engineering problems, data and applications are the books by Gnedenko, Belyaev, and Soloyev (1969) Mann, Schafer, and Singpurwalla (1974), Crowder *et al.* (1991) and Meeker and Escobar (1998); sandwiched between these are a myriad of others (cf. Singpurwalla 1993). I have mentioned these three books because the first is one of the oldest and the third, one of the latest. In the area of survival analysis, a classic source is the book by Kalbfleisch and Prentice (1980). Recent additions to this field include several books on counting process models, the one by Anderson *et al.* (1993) being encyclopedic. In all the above books, the statistical paradigm that is subscribed to is different from what we propose to do here. Thus, the possibilities of duplication appear to be minimum. As a final comment, in Barlow, Clarotti and Spizzichino (1993), the need to integrate reliability analysis into risk management has been recognized, but the material is more along the lines of a research monograph rather than an expository development.

1.6 OVERVIEW OF THE BOOK

Chapters 2 and 3 pertain to foundational issues; they present the underlying paradigm for our development. Chapter 2 starts off with a discussion of uncertainty and its quantification, leading up to decision making under uncertainty. In the interim we also present standard statistical notions such as ‘inference’, ‘likelihood’ and ‘prediction’. Professionally trained statisticians, biostatisticians and applied probabilists may find little, if any, that is new to them here. By and large, the material of Chapter 2 is pitched toward engineers, operations research analysts and other mathematically oriented subject matter specialists. The material may also appeal to graduate students in the statistical sciences, and other statisticians who may not have had an exposure to subjective Bayesian thinking. Chapter 3 is specialized and pertains to the important notion of ‘exchangeability’, which plays a key role in selecting models for quantifying uncertainty. Most readers may want to skip this material on a first reading. Chapter 4 is basic and pertains to a discussion of standard notions in reliability and survival analysis. However, the perspective that we take in Chapter 4 is not traditional; the material given here should be viewed as being foundational to reliability and survival analysis. Chapter 5 presents a different perspective on the same topics that are covered in the books of category three mentioned before. Chapter 6 builds on the material of Chapter 5; it pertains to the propagation of uncertainty through a system of items. The material in the remaining chapters is mildly advanced and pertains to specialized topics,

8 INTRODUCTION AND OVERVIEW

many of which may appeal to only certain segments of the readership. Thus, Chapter 7 focuses on dynamic environments and the use of stochastic process models, Chapter 8 on counting processes and event history data and Chapter 9 on non-parametric methods within the Bayesian paradigm. Chapter 10 pertains to the survivability of systems with interdependent failures and Chapter 11 describes the role of reliability and survival analysis in econometrics, asset pricing and mathematical finance. Our overall aim has been to give as broad a coverage as is possible, even if this means an occasional sacrifice of specifics. We compensate for this compromise of completeness by providing adequate references so that a reader can patch together a more complete picture.