

Cambridge University Press
978-0-521-77320-1 - Probabilistic Risk Analysis: Foundations and Methods
Tim Bedford and Roger Cooke
Excerpt
[More information](#)

Part I

Introduction

Cambridge University Press
978-0-521-77320-1 - Probabilistic Risk Analysis: Foundations and Methods
Tim Bedford and Roger Cooke
Excerpt
[More information](#)

1

Probabilistic risk analysis

Probabilistic risk analysis (PRA), also called quantitative risk analysis (QRA) or probabilistic safety analysis (PSA), is currently being widely applied to many sectors, including transport, construction, energy, chemical processing, aerospace, the military, and even to project planning and financial management. In many of these areas PRA techniques have been adopted as part of the regulatory framework by relevant authorities. In other areas the analytic PRA methodology is increasingly applied to validate claims for safety or to demonstrate the need for further improvement. The trend in all areas is for PRA to support tools for management decision making, forming the new area of *risk management*.

Since PRA tools are becoming ever more widely applied, and are growing in sophistication, one of the aims of this book is to introduce the reader to the main tools used in PRA, and in particular to some of the more recent developments in PRA modeling. Another important aim, though, is to give the reader a good understanding of uncertainty and the extent to which it can be modeled mathematically by using probability. We believe that it is of critical importance not just to understand the mechanics of the techniques involved in PRA, but also to understand the foundations of the subject in order to judge the limitations of the various techniques available. The most important part of the foundations is the study of uncertainty. What do we mean by uncertainty? How might we quantify it?

After the current introductory chapter, in Part two we discuss theoretical issues such as the notion of uncertainty and the basic tools of probability and statistics that are widely used in PRA. Part three presents basic modeling tools for engineering systems, and discusses some of the techniques available to quantify uncertainties (both on the basis of reliability data, and using expert judgement). In Part four we discuss uncertainty modeling and risk measurement. The aim is to show how dependent uncertainties are important

and can be modeled, how value judgements can be combined with uncertainties to make optimal decisions under uncertainty, and how uncertainties and risks can be presented and measured.

1.1 Historical overview

This introductory section reviews the recent history of these developments, focusing in particular on the aerospace, nuclear and chemical process sectors.

1.1.1 The aerospace sector

A systematic concern with risk assessment methodology began in the aerospace sector following the fire of the Apollo test AS-204 on January 27, 1967, in which three astronauts were killed. This one event set the National Aeronautics and Space Administration (NASA) back 18 months, involved considerable loss of public support, cost NASA salaries and expenses for 1500 people involved in the subsequent investigation, and ran up \$410 million in additional costs [Wiggins, 1985]. Prior to the Apollo accident, NASA relied on its contractors to apply ‘good engineering practices’ to provide quality assurance and quality control.

On April 5, 1969 the Space Shuttle Task Group was formed in the Office of Manned Space Flight of NASA. The task group developed ‘suggested criteria’ for evaluating the safety policy of the shuttle program which contained quantitative safety goals. The probability of mission completion was to be at least 95% and the probability of death or injury per mission was not to exceed 1%. These numerical safety goals were not adopted in the subsequent shuttle program.

The reason for rejecting quantitative safety goals given at the time was that low numerical assessments of accident probability could not guarantee safety: ‘... the problem with quantifying risk assessment is that when managers are given numbers, the numbers are treated as absolute judgments, regardless of warnings against doing so. These numbers are then taken as fact, instead of what they really are: subjective evaluations of hazard level and probability.’ ([Wiggins, 1985], p. 85).

An extensive review of the NASA safety policy following the Challenger accident of January 28, 1986 brought many interesting facts to light. A quantitative risk study commissioned by the US Air Force in 1983 estimated the Challenger’s solid rocket booster failure probability per launch as 1 in 35. NASA management rejected this estimate and elected to rely on

their own engineering judgment, which led to a figure of 1 in 100,000 [Colglazier and Weatherwas, 1986]

It has also become clear that distrust of reassuring risk numbers was not the reason that quantitative risk assessment was abandoned. Rather, initial estimates of catastrophic failure probabilities were so high that their publication would have threatened the political viability of the entire space program. For example, a General Electric ‘full numerical probabilistic risk assessment’ on the likelihood of successfully landing a man on the moon indicated that the chance of *success* was ‘less than 5%’. When the NASA administrator was shown these results, he ‘felt that the numbers could do irreparable harm, and disbanded the effort’ [Bell and Esch, 1989].

By contrast, a congressional report on the causes of the Shuttle accident (quoted in [Garrick, 1989]) concluded that the qualitative method of simply identifying failures leading to loss of vehicle accidents (the so-called *critical items*) was limited because not all elements posed an equal threat. Without a means of identifying the *probability* of failure of the various elements NASA could not focus its attention and resources effectively.

Since the shuttle accident, NASA has instituted programs of quantitative risk analysis to support safety during the design and operations phases of manned space travel. The NASA risk assessment effort reached a high point with the publication of the SAIC Shuttle Risk Assessment [Fragola, 1995]. With this assessment in hand, NASA was able to convince the US Congress that the money spent on shuttle development since the Challenger accident had been well used, even though no failure paths had been eliminated. The report showed that the *probability* of the most likely failure causes had been significantly reduced.

The European space program received a setback with the failure of the maiden flight of Ariane 5. A board of inquiry [ESA, 1997] revealed that the disaster was caused by software errors and the management of the software design. The accident demonstrates the problem of integrating working technologies from different environments into a new reliable system.

1.1.2 The nuclear sector

Throughout the 1950s, following Eisenhower’s ‘Atoms for Peace’ program, the American Atomic Energy Commission (AEC) pursued a philosophy of risk assessment based on the ‘maximum credible accident’. Because ‘credible accidents’ were covered by plant design, residual risk was estimated by studying the hypothetical consequences of ‘incredible accidents’. An early study released in 1957 focused on three scenarios of radioactive releases

from a 200 megawatt nuclear power plant operating 30 miles from a large population center. Regarding the probability of such releases the study concluded that ‘no one knows now or will ever know the exact magnitude of this low probability’.

Successive design improvements were intended to reduce the probability of a catastrophic release of the reactor core inventory. Such improvements could have no visible impact on the risk as studied with the above methods. On the other hand, plans were being drawn for reactors in the 1000 megawatt range located close to population centers, and these developments would certainly have a negative impact on the consequences of the ‘incredible accident’.

The desire to quantify and evaluate the effects of these improvements led to the introduction of *probabilistic* risk analysis. As mentioned above, the basic methods of probabilistic risk assessment originated in the aerospace program in the 1960s. The first full scale application of these methods, including an extensive analysis of the accident consequences, was undertaken in the *Reactor Safety Study* WASH-1400 [NRC, 1975] published by the US Nuclear Regulatory Commission (NRC). This study is rightly considered to be the first modern PRA.

The reception of the *Reactor Safety Study* in the scientific community may best be described as turbulent. The American Physical Society [APS, 1975] conducted an extensive review of the first draft of the *Reactor Safety Study*. In the accompanying letter attached to their report, physicists Wolfgang Panofsky, Victor Weisskopf and Hans Bethe concluded, among other things, that the calculation methods were ‘fairly unsatisfactory’, that the emergency core cooling system is unpredictable and that relevant physical processes ‘which could interfere with its functioning have not been adequately analyzed’, and that ‘the consequences of an accident involving major radioactive release have been underestimated as to casualties by an order of magnitude’. The final draft of the *Reactor Safety Study* was extensively reviewed by, among others, the Environmental Protection Agency [EPA, 1976] and the Union of Concerned Scientists [Union of Concerned Scientists, 1977].

In 1977 the United States Congress passed a bill creating a special ‘review panel’ of external reactor safety experts to review the ‘achievements and limitations’ of the *Reactor Safety Study*. The panel was led by Prof. Harold Lewis, and their report is known as the ‘Lewis Report’ [Lewis *et al.*, 1979]. While the Lewis Report recognized the basic validity of the PRA methodology and expressed appreciation for the pioneering effort put into the *Reactor Safety Study*, they also uncovered many deficiencies in the treatment of probabilities. They were led to conclude that the uncertainty bands claimed for the conclusions in the *Reactor Safety Study* were ‘greatly understated’.

Significantly, the Lewis Report specifically endorsed the use of subjective probabilities in the *Reactor Safety Study*.

In January 1979 the NRC distanced itself from the results of the *Reactor Safety Study*: ‘In particular, in light of the Review Group conclusions on accident probabilities, the Commission does not regard as reliable the *Reactor Safety Study*’s numerical estimate of the overall risk of a reactor accident.’ [NRC, 1979]

The future of PRA after the NRC’s announcement of 1979 did not look bright. However, the dramatic events of March 1979 served to change that. In March 1979 the Three Mile Island – 2 (TMI) Nuclear Generating Unit suffered a severe core damage accident. Subsequent study of the accident revealed that the accident sequence had been predicted by the *Reactor Safety Study*. The probabilities associated with that sequence, particularly those concerning human error, do not appear realistic in hindsight.

Two influential independent analyses of the TMI accident, the Report of the President’s Commission on the Accident at Three Mile Island [Kemeny *et al*, 1979] and the Rogovin Report [Rogovin and Frampton, 1980], credited the *Reactor Safety Study* with identifying the small loss-of-coolant accidents as the major threat to safety, and recommended that greater use should be made of probabilistic analyses in assessing nuclear plant risks. They also questioned whether the NRC was capable of regulating the risks of nuclear energy, and recommended that the regulatory body be massively overhauled (which recommendation was not carried out).

Shortly thereafter a new generation of PRAs appeared in which some of the methodological defects of the *Reactor Safety Study* were avoided. The US NRC released *The Fault Tree Handbook* [Vesely *et al.*, 1981] in 1981 and the *PRA Procedures Guide* [NRC, 1983] in 1983 which shored up and standardized much of the risk assessment methodology. Garrick’s review [Garrick, 1984] of PRAs conducted in the aftermath of the Lewis report discussed the main contributors to accident probability identified at the plants. He also noted the necessity to model uncertainties properly in order to use PRA as a management tool, and suggested the use of on-line computer PRA models to guide plant management (a process now called a ‘living PRA’).

The accident at the Chernobyl Nuclear Power Plant occurred on April 26, 1986. A test was carried out in order to determine how long the reactor coolant pumps could be operated using electrical power from the reactors’ own turbine generator under certain abnormal conditions. At the beginning of the test some of the main coolant pumps slowed down, causing a reduction of coolant in the core. The coolant left began to boil, adding reactivity to

the core (due to the so-called positive void coefficient of the RBMK reactor type). This caused a sudden increase in power which could not be controlled because the control systems worked too slowly. The power surge caused the fuel to overheat and disintegrate. Pieces of fuel ejected into the coolant then caused a steam explosion whose force blew the cover off the reactor. There were 31 early deaths and (amongst other radiological effects) a ‘real and significant’ increase in childhood carcinoma of the thyroid [OECD, 1996]. Thousands of people have been displaced and blame the reactor accident for all sorts of health problems.

The management of western nuclear power corporations moved quickly to assure the public that this type of accident could not occur in the US and western Europe because of the difference in reactor design. The Three Mile Island and Chernobyl accidents, in addition to regular publicity about minor leaks of radioactive material from other power stations and processing plants, however, have fostered a climate of distrust in nuclear power and in the capacity of management to run power stations properly.

Besides technical advances in the methodology of risk analyses, the 1980s and 1990s have seen the further development of numerical safety goals. Examples are the USNRC policy statement of 1986 [NRC, 1986] and the UK *Tolerability of risk* document [HSE, 1987]. These documents seek to place the ALARP principle (‘as low as reasonably possible’) into a numerical framework by defining upper levels of intolerable risk and lower levels of broadly tolerable risk.

1.1.3 The chemical process sector

In the process sector, government authorities became interested in the use of probabilistic risk analysis as a tool for estimating public exposure to risk in the context of licensing and citing decisions. Important European efforts in this direction include two studies of refineries on Canvey Island ([HSE, 1978], [HSE, 1981]) in the UK, a German sulphuric acid plant study [Jäger, 1983], and the Dutch LPG and COVO studies ([TNO, 1983], [COVO, 1982]). The COVO study [COVO, 1982] was a risk analysis of six potentially hazardous objects in the Rijnmond area. The group which performed the study later formed the consulting firm Technica, which has since played a leading role in risk analysis.

The impetus for much of this work was the Post-Seveso Directive [EEC, 1982] adopted by the European Community following the accidental release of dioxin by a chemical plant near Seveso, Italy. The directive institutes a policy of risk management for chemical industries handling hazardous

materials in which each member state has the responsibility for developing its own risk management methodology. The Dutch government took the lead in requiring quantitative risk analyses of potentially hazardous objects, and has invested heavily in developing tools and methods for implementing this policy. This lead has been followed by several other member countries.

Dutch legislation requires the operator of a facility dealing with hazardous substances to submit an 'External Safety Report' (Externe Veiligheid Rapport, EVR). About 70 companies in the Netherlands fall under this reporting requirement. EVRs are to be updated every 5 years.

The quantitative risk analysis required in EVRs may be broken down into four parts. The first part identifies the undesirable events which may lead to a threat to the general public. The second part consists of an effect and damage assessment of the undesired events. The third part calculates the probability of damage, consisting of the probability of an undesired release of hazardous substances, and the probability of propagation through the environment causing death. The fourth part determines the individual and group risk associated with the installation.

For new installations, the risk of death of an 'average individual' exposed at any point outside the installation perimeter for an entire year must not exceed 10^{-6} . The group risk requirement stipulates that the probability of 10^n or more fatalities ($n > 1$) must not exceed 10^{-3-2n} per year. If this probability exceeds 10^{-5-2n} per year then further reduction is required.

Risk based regulation is now common in many different sectors. An overview of some of the risk goals currently set is given in Table 18.5.

1.1.4 The less recent past

We have concentrated on the developments in a few important sectors since the end of the second world war. The reader interested in looking further into the past is referred to [Covello and Mumpower, 1985] which starts about 3200 BC, and [Bernstein, 1996].

1.2 What is the definition of risk?

The literature on the subject of risk has grown rapidly in the last few years, and the word 'risk' is used in many different ways. The purpose of this section is to discuss briefly what we mean by risk, and in what way the concept can be described in a mathematical setting. The limitations of the mathematical approach to measuring risk will be highlighted in Chapter 18. Our discussion here is largely drawn from [Kaplan and Garrick, 1981].




Fig. 1.1. Risk curve

A *hazard* is considered as a source of danger but the concept does not contain any notion of the likelihood with which that danger will actually impact on people or on the environment. We are often *uncertain* about whether or not a hazard will actually lead to negative consequences (that is, whether the potentiality will be converted into actuality). As will be argued in Chapter 2, that uncertainty can – in principle – be quantified by *probability*. The definition of *risk* combines both of the above elements. A risk analysis tries to answer the questions:

- (i) What can happen?
- (ii) How likely is it to happen?
- (iii) Given that it occurs, what are the consequences?

Kaplan and Garrick [Kaplan and Garrick, 1981] define risk to be a set of *scenarios* s_i , each of which has a *probability* p_i and a *consequence* x_i . If the scenarios are ordered in terms of increasing severity of the consequences then a risk curve can be plotted, for example as shown in Figure 1.1. The risk curve illustrates what is the probability of *at least* a certain number of casualties in a given year.

Kaplan and Garrick [Kaplan and Garrick, 1981] further refine the notion of risk in the following way. First, instead of talking about the probability of an event, they talk about the *frequency* with which such an event might take place. They then introduce the notion of uncertainty about the frequency (the ‘probability of a frequency’). This more sophisticated notion of risk will be discussed further in Chapter 2.