1

Introduction

Risk assessment is a methodological framework for determining the nature and extent of the risk associated with an activity. It comprises the following three main steps:

- Identification of relevant sources of risk (threats, hazards, opportunities)
- Cause and consequence analysis, including assessments of exposures and vulnerabilities
- Risk description.

Risk assessment is now widely used in the context of various types of activities as a tool to support decision making in the selection of appropriate protective and mitigating arrangements and measures, as well as in ensuring compliance with requirements set by, for example, regulatory agencies. The basis of risk assessment is the systematic use of analytical methods whose quantification is largely probability based. Common methods used to systematically analyze the causes and consequences of failure configurations and accident scenarios are fault trees and event trees, Markov models, and Bayesian belief networks; statistical methods are used to process the numerical data and make inferences. These modeling methods have been developed to gain knowledge about cause–effect relationships, express the strength of these relationships, characterize the remaining uncertainties, and describe, in quantitative or qualitative form, other properties relevant for risk management (IAEA, 1995; IEC, 1993). In short, risk assessments specify what is at stake, assess the uncertainties of relevant quantities, and produce a risk description which provides information useful for the decision-making process of risk management.

In this book we put the main focus on quantitative risk assessment (QRA), where risk is expressed using an adequate representation of the uncertainties involved. To further develop the methodological framework of risk assessment, we will need to explain in more detail what we mean by risk.

Uncertainty in Risk Assessment: The Representation and Treatment of Uncertainties by Probabilistic and Non-Probabilistic Methods, First Edition. Terje Aven, Piero Baraldi, Roger Flage and Enrico Zio.

This introductory chapter is organized as follows. Following Section 1.1, which addresses the risk concept, we present in Section 1.2 the main features of probabilistic risk assessment (PRA), which is a QRA based on the use of probability to characterize and represent the uncertainties. Then, in Section 1.3, we discuss the use of risk assessment in decision-making contexts. Section 1.4 considers the issue of uncertainties in risk assessment, motivated by the thesis that if uncertainty cannot be properly treated in risk assessment, the risk assessment tool fails to perform as intended (Aven and Zio, 2011). This section is followed by a discussion on the main challenges of the probability-based approaches to risk assessment, and the associated uncertainty analysis. Alternative approaches for dealing with uncertainty are briefly discussed.

1.1 Risk

1.1.1 The concept of risk

In all generality, risk arises wherever there exists a potential source of damage or loss, that is, a hazard (threat), to a target, for example, people, industrial assets, or the environment. Under these conditions, safeguards are typically devised to prevent the occurrence of the hazardous conditions, and protection is put in place to counter and mitigate the associated undesired consequences. The presence of a hazard does not in itself suffice to define a condition of risk; indeed, inherent in the latter there is the uncertainty that the hazard translates from potential to actual damage, bypassing safeguards and protection. In synthesis, the notion of risk involves some kind of loss or damage that might be received by a target and the uncertainty of its transformation in actual loss or damage, see Figure 1.1. Schematically we can write (Kaplan and Garrick, 1981; Zio, 2007; Aven, 2012b)

$$\text{Risk} = \text{Hazards/Threats and Consequences (damage)} + \text{Uncertainty}. \qquad (1.1)$$

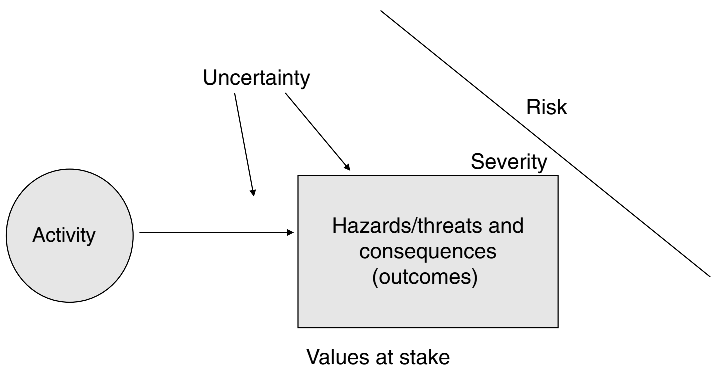

Figure 1.1 The concept of risk reflecting hazards/threats and consequences and associated uncertainties (what events will occur, and what the consequences will be).

Normally, the consequence dimension relates to some type of undesirable outcome (damage, loss, harm). Note that by centering the risk definition around undesirable outcomes, we need to define what is undesirable, and for whom. An outcome could be positive for some stakeholders and negative for others: discussing whether an outcome is classified in the right category may not be worth the effort, and most of the general definitions of risk today allow for both positive and negative outcomes (Aven and Renn, 2009).

Let A denote a hazard/threat, C the associated consequences, and U the uncertainties (will A occur, and what will C be?). The consequences relate to something that humans value (health, the environment, assets, etc.). Using these symbols we can write (1.1) as

$$\text{Risk} = (A, C, U), \tag{1.2}$$

or simply

$$\text{Risk} = (C, U), \tag{1.3}$$

where C in (C, U) expresses all consequences of the given activity, including the hazardous/threatful events A. These two risk representations are shown in Figure 1.2.

Obviously, the concept of risk cannot be limited to one particular measuring device (e.g., probability) if we seek a general risk concept. For the measure introduced, we have to explain precisely what it actually expresses. We also have to clarify the limitations with respect to its ability to measure the uncertainties: is there a need for a supplement to fully describe the risk? We will thoroughly discuss these issues throughout the book.

A concept closely related to risk is vulnerability (given the occurrence of an event A). Conceptually vulnerability is the same as risk, but conditional on the occurrence of an event A:

$$\text{Vulnerability} \mid A = \text{Consequences} + \text{Uncertainty} \mid \text{the occurrence of the event } A, \tag{1.4}$$

where the symbol $\mid$ indicates "given" or "conditional." For short we write

$$\text{Vulnerability} \mid A = (C, U \mid A). \tag{1.5}$$

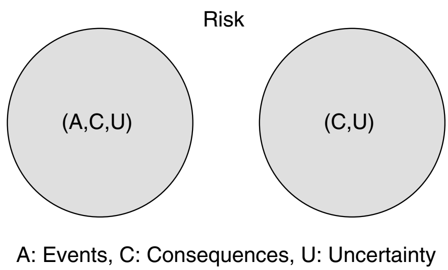

Figure 1.2 The main components of the concept of risk used in this book.

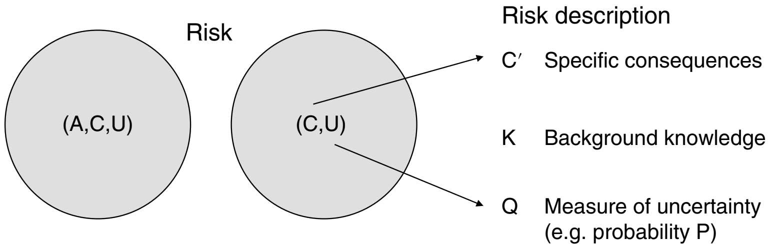

Figure 1.3 Illustration of how the risk description is derived from the concept of risk.

1.1.2 Describing/measuring risk

The risk concept has been defined above. However, this concept does not give us a tool for assessing and managing risk. For this purpose we must have a way of describing or measuring risk, and the issue is how.

As we have seen, risk has two main dimensions, consequences and uncertainty, and a risk description is obtained by specifying the consequences C and using a description (measure) of the uncertainty, Q. The most common tool is probability P, but others exist and these also will be given due attention in the book. Specifying the consequences means identifying a set of quantities of interest C' that represent the consequences C, for example, the number of fatalities.

Now, depending on the principles laid down for specifying C' and the choice of Q, we obtain different perspectives on how to describe/measure risk. As a general description of risk we can write

$$\begin{aligned} &\text{Risk description} = (C', Q, K), \\ &(\text{or, alternatively}, (A', C', Q, K)), \end{aligned} \tag{1.6}$$

where K is the background knowledge (models and data used, assumptions made, etc.) that Q and the specification C' are based on, see Figure 1.3. On the basis of the relation between vulnerability and risk previously introduced, the vulnerability given an event A is analogously described by $(C', Q, K|A)$.

1.1.3 Examples

1.1.3.1 Offshore oil and gas installation

Consider the future operation of an offshore installation for oil and gas processing. We all agree that there is some "risk" associated with this operation. For example, fires and explosions could occur leading to fatalities, oil spills, economic losses, and so on. Today we do not know if these events will occur and what the specific consequences will be: we are faced with uncertainties and, thus, risk. Risk is two dimensional, comprising events and consequences, and associated uncertainties (i.e., the events and consequences being unknown, the occurrences of the events are not known and the consequences are not known).

When performing a risk assessment we describe and/or quantify risk, that is, we specify (C', Q, K). For this purpose we need quantities representing C' and a measure of uncertainty; for the latter, probability is introduced. Then, in the example discussed, C' is represented by the number of fatalities, $Q=P$, and the background knowledge K covers a number of assumptions that the assessment is based on, for example, related to the number of people working on the installation, as well as the models and data used for quantification of the accident probabilities and consequences. On this basis, several risk indices or metrics are defined, such as the expected number of fatalities (e.g., potential loss of lives, PLL, typically defined for a one-year period) and the fatal accident rate (FAR, associated with 100 million exposed hours), the probability that a specific person will be killed in an accident (individual risk, IR), and frequency–consequence (f–n) curves expressing the expected number of accidents (frequency f) with at least n fatalities.

1.1.3.2 Health risk

Consider a person's life and focus on the condition of his/her health. Suppose that the person is 40 years old and we are concerned about the "health risk" for this person for a predetermined period of time or for the rest of his/her life. The consequences of interest in this case arise from "scenarios" of possible specific diseases (known or unknown types) and other illnesses, their times of development, and their effects on the person (will he/she die, suffer, etc.).

To describe risk in this case we introduce the frequentist probability p that the person gets a specific disease (interpreted as the fraction of persons that get the disease in an infinite population of "similar persons"), and use data from a sample of "similar persons" to infer an estimate p^* of p. The probability p can be considered a parameter of a binomial probability model.

For the consequent characterization, C', we look at the occurrence or not of a disease for the specific person considered, and the time of occurrence of the disease, if it occurs. In addition, we have introduced a probability model with a parameter p and this p also should be viewed as a quantity of interest C'. We seek to determine p, but there are uncertainties about p and we may use confidence intervals to describe this uncertainty, that is, to describe the stochastic variation in the data.

The uncertainty measure in this case is limited to frequentist probabilities. It is based on a traditional statistical approach. Alternatively, we could have used a Bayesian analysis based on subjective (judgmental, knowledge-based) probabilities P (we will return to the meaning of these probabilities in Chapter 2). The uncertainty description in this case may include a probability distribution of p, for example, expressed by the cumulative distribution function $F(p')=P(p \le p')$. Using P to measure the uncertainties (i.e., $Q=P$), we obtain a risk description (C', P, K), where p is a part of C'. From the distribution $F(p')$ we can derive the unconditional probability $P(A)$ (more precisely, $P(A|K)$) of the event A that the person gets the disease, by conditioning on the true value of p (see also Section 2.4):

$$P(A) = \int P(A \mid p')\, \mathrm{d}F(p') = \int p' \, \mathrm{d}F(p'). \tag{1.7}$$

This probability is a subjective probability, based on the probability distribution of the frequentist probability p. We see that $P(A)$ is given by the center of gravity (the expected value) of the distribution F.

Alternatively, we could have made a direct subjective probability assignment for $P(A) = P(A \mid K)$, without introducing the probability model and the parameter p.

1.2 Probabilistic risk assessment

Since the mid-1970s, the framework of probability theory has been the basis for the analytic process of risk assessment (NRC, 1975); see the reviews by Rechard (1999, 2000). A probabilistic risk assessment (PRA) systematizes the knowledge and uncertainties about the phenomena studied: what are the possible hazards and threats, their causes and consequences? The knowledge and uncertainties are characterized and described using various probability-based metrics, as illustrated in Section 1.1.3; see also Jonkman, van Gelder, and Vrijling (2003) for a comprehensive overview of risk metrics (indices) for loss of life and economic damage. Additional examples will be provided in Chapter 3, in association with some of the detailed modeling and tools typical of PRA.

A total PRA for a system comprises the following stages:

1. *Identification of threats/hazards.* As a basis for this activity an analysis of the system is carried out in order to understand how the system works, so that departures from normal, successful operation can be identified. A first list of hazards/threats is normally identified based on this system analysis, as well as on experience from similar types of analyses, statistics, brainstorming activities, and specific tools such as failure mode and effect analysis (FMEA) and hazards and operability (HAZOP) studies.

2. *Cause analysis.* In cause analysis, we study the system to identify the conditions needed for the hazards/threats to occur. What are the causal factors? Several techniques exist for this purpose, from brainstorming sessions to the use of fault tree analyses and Bayesian networks.

3. *Consequence analysis.* For each identified hazard/threat, an analysis is carried out addressing the possible consequences the event can lead to. Consequence analysis deals to a large extent with the understanding of physical phenomena, for example, fires and explosions, and various types of models of the phenomena are used. These models may for instance be used for answering questions like: How will a fire develop? What will be the heat at various distances? What will the explosive pressure be in case an explosion takes place? And so on. Event tree analysis is a common method for analyzing the

scenarios that can develop in the different consequences. The number of steps in the sequence of events that form a scenario is mainly dependent on the number of protective barriers set up in the system to counteract the initiating event of that sequence. The aim of the consequence-reducing barriers is to prevent the initiating events from resulting in serious consequences. For each of these barriers, we can carry out failure analysis to study their reliability and effectiveness. Fault tree analysis is a technique often used for this purpose.

4. *Probabilistic analysis.* The previous stages of analysis provide a set of sequences of events (scenarios), which lead to different consequences. This specification of scenarios does not address the question of how likely the different scenarios and the associated consequences are. Some scenarios could be very serious, should they occur, but if the likelihood of their occurrence is low, they are not so critical. Using probability models to reflect variation in the phenomena studied and assigning probabilities for the occurrence of the various events identified and analyzed in steps 2 and 3, overall probability values and expected consequence values can be computed.

5. *Risk description.* Based on the cause analysis, consequence analysis, and probabilistic analysis, risk descriptions can be obtained using various metrics, for example, risk matrices showing the computed/assigned probability of a hazard/threat and the expected consequences given that this event has occurred, as well as IR, PLL, and FAR values.

6. *Risk evaluation.* The results of the risk analysis are compared to predefined criteria, for example, risk tolerability limits or risk acceptance criteria.

PRA methodology is nowadays used extensively in industries such as nuclear power generation (e.g., Vesely and Apostolakis, 1999; Apostolakis, 2004), offshore petroleum activities (e.g., Falck, Skramstad, and Berg, 2000; Vinnem, 2007), and air transport (e.g., Netjasov and Janic, 2008).

The current default approach to a comprehensive quantitative PRA is based on the so-called set of triplets definition of risk, introduced by Kaplan and Garrick (1981); see also Kaplan (1992, 1997). In this approach, risk is defined as the combination of possible scenarios s, resulting consequences x, and the associated likelihoods l. Loosely speaking: What can happen (go wrong)? How likely is it? What are the consequences? Within this conceptual framework, three main likelihood settings are often defined (Kaplan, 1997): repetitive situation with known frequency ($l = f$, where f is a known frequentist probability), unique situation ($l = p$, where p is a subjective probability), and repetitive situation with unknown frequency ($l = H(f)$, where H is a subjective probability distribution on an unknown/uncertain frequentist probability f). Of course, the first case is a special case of the third. The last-mentioned setting is typically dealt with using the so-called probability of frequency approach, where all potentially occurring events involved are assumed to have uncertain frequency probabilities of occurrence, and the epistemic uncertainties about the true values of frequency probabilities are described using subjective probabilities. For the sake of simplicity, in the following we will often use the short term “frequency” instead of “frequentist probability.”

The probability of frequency approach is in line with the standard Bayesian approach (Aven, 2012a) as will be described below. It is also considered “the most general and by far the most powerful and useful idea” by Kaplan (1997, p. 409), and corresponds to the highest level of sophistication in the treatment of uncertainties in risk analysis according to the classification by Paté-Cornell (1996).

In this book, however, we adopt a broader perspective of risk by which the set of triplets is not *risk* per se but a *risk description*. In this view, the outcome of a risk assessment is a list of scenarios quantified in terms of probabilities and consequences, which collectively describe the risk. As we will thoroughly discuss throughout the book, this risk description will be shown to be more or less adequate for describing the risk and uncertainties in different situations.

Numerous textbooks deal with methods and models for PRA, for example, Andrews and Moss (2002), Aven (2008), Cox (2002), Vinnem (2007), Vose (2008), and Zio (2007, 2009). Some also deal specifically with foundational issues, in particular with the concepts of uncertainty and probability, for example, Aven (2012a), Bedford and Cooke (2001), and Singpurwalla (2006).

In spite of the maturity reached by the methodologies used in PRA, a number of new and improved methods have been developed in recent years to meet the needs of the analysis brought about by the increasing complexity of the systems and processes studied, and to respond to the introduction of new technological systems. Many of the methods introduced allow for increased levels of detail and precision in the modeling of phenomena and processes within an integrated framework of analysis covering physical phenomena, human and organizational factors, and software dynamics (e.g., Mohaghegh, Kazemi, and Mosleh, 2009). Other methods are devoted to the improved representation and assessment of risk and uncertainty. Examples of more recently developed methods are Bayesian belief networks, binary digit diagrams, multi-state reliability analysis, and advanced Monte Carlo simulation tools. For a summary and discussion of some of these models and techniques, see Bedford and Cooke (2001) and Zio (2009).

The probabilistic analysis underpinning PRA is based on one or the other of two alternative conceptual foundations: the traditional frequentist approach and the Bayesian approach (Bedford and Cooke, 2001; Aven, 2012a). The former is typically applied in situations in which there exists a large amount of relevant data; it is founded on well-known principles of statistical inference, the use of probability models, the interpretation of probabilities as relative frequencies, point estimates, confidence interval estimation, and hypothesis testing.

By contrast, the Bayesian approach is based on the concept of subjective (judgmental, knowledge-based) probabilities and is applied in situations in which there exists only a limited amount of data (e.g., Guikema and Paté-Cornell, 2004). The idea is to first establish probability models that adequately represent the aleatory uncertainties, that is, the inherent variability of the phenomena studied, such as the distribution of lifetimes of a type of system. The epistemic uncertainties, reflecting incomplete knowledge or lack of knowledge about the values of the parameters of the models, are then represented by prior subjective probability distributions. When new data on the phenomena studied becomes available, Bayes' formula is used to update the representation of the epistemic uncertainties in terms of the posterior distributions. Finally, the predictive distributions of the quantities of interest – the observables (e.g., the lifetime of new systems) – are derived by applying the law of total probability. The predictive distributions are epistemic statements, but they also reflect the inherent variability of the phenomena being studied, that is, the aleatory uncertainties.

1.3 Use of risk assessment: The risk management and decision-making context

Risk management can be defined as the coordinated activities to direct and control an organization with regard to risk (ISO, 2009). As illustrated in Figure 1.4, the main central steps of the risk management process are: establishment of the context, risk assessment, and risk treatment. Context here refer to the internal and external environment of the organization, the interface of these environments, the purpose of the risk management activity, and suitable risk criteria. Risk treatment is the process of modifying risk, which may involve avoiding, modifying, sharing or retaining risk (ISO, 2009).

Note that, according to ISO (2009), source (hazard/threat/opportunity) identification is not included as part of risk analysis. Many analysts and researchers do prefer, however, to include this element in the definition of the scope of risk analysis, in

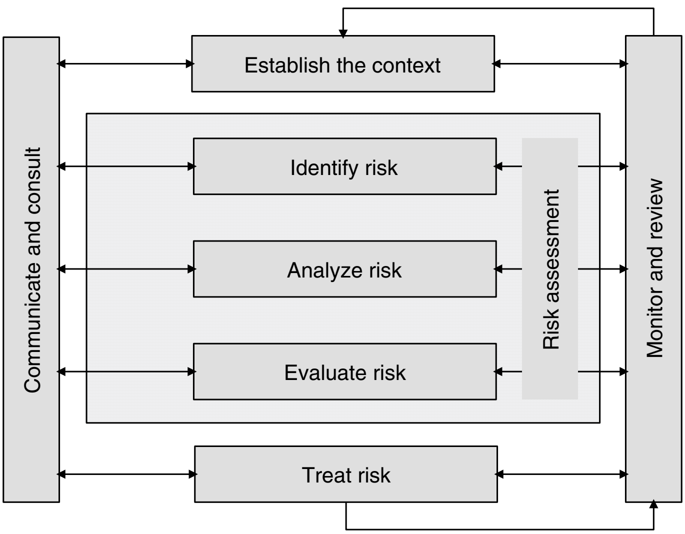

Figure 1.4 The risk management process (based on ISO, 2009).

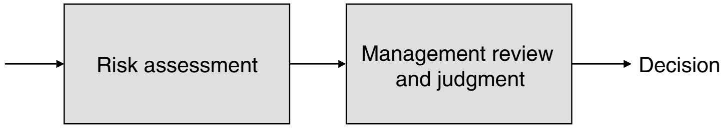

Figure 1.5 The leap (the management review and judgment) between risk assessment and the decision.

addition to cause and consequence analysis and risk description; see, for example, Modarres, Kamiskiy, and Krivtsov (1999) and Aven (2008).

There are different perspectives on how to *use* risk assessment and uncertainty analysis in risk management and decision making. Strict adherence to expected utility theory, cost–benefit analysis, stochastic optimization, and related theories would give clear recommendations on what is the optimal arrangement or measure. However, most risk researchers and risk analysts would see risk and uncertainty assessments as decision support tools, in the sense that the assessments inform the decision makers. The decision making is risk informed, not risk based (Apostolakis, 2004). In general, there is a significant leap from the assessments to the decision, see Figure 1.5. What this leap (often referred to as management review and judgment, or risk evaluation) comprises is a subject being discussed in the literature (e.g., Renn, 2005; Aven and Renn, 2010; Aven and Vinnem, 2007). The management review and judgment is about giving weight to the cautionary and precautionary policies and risk perception, as well as other concerns/attributes than risk and uncertainties. The scope and boundaries of risk and uncertainty assessments define to a large extent the content of the review and judgment, as we will discuss in Section 1.4.

Similar ideas are reflected in many risk assessment frameworks, for example, the analytic–deliberative process recommended by the US National Research Council (1996, 2008) in environmental restoration decisions involving multiple stakeholders.

The relevance of the management and decision-making side of the problem is further demonstrated by the increasing research efforts being conducted to integrate decision-making and uncertainty characterizations that extend beyond the traditional probability-based representations. An example is the theory of robust decision making, supported by fairly recent advancements in robust optimization (Ben-Tal and Nemirovski, 2002; Beyes and Sendhoff, 2007), aimed at finding optimal decisions under specified ranges of the uncertain model parameters.

In this book we will present and discuss various ways of representing and characterizing the uncertainties in risk assessment. The key question addressed is how to best express risk and represent the associated uncertainties to meet the decision makers' and other relevant stakeholders' needs, in the typical settings of risk assessment of complex systems with limited knowledge of the behavior of these systems. The principal driver is the decision-making process and the need to inform and facilitate this process with representative information derived from the risk assessment.

Hence, the book addresses, for example, the issue of how to present the results from risk and uncertainty assessments to decision makers, but not the decision making itself. We make a sharp distinction between risk/uncertainty representation and characterization on the one hand, and risk/uncertainty management and related decision making on the other.

1.4 Treatment of uncertainties in risk assessments

When speaking about uncertainties in risk assessments most analysts would think about the uncertainties related to parameters in probability models, such as the frequentist probability p in the second example in Section 1.1.3. Following the traditional statistical approach, the uncertainties are expressed using confidence intervals or according to the Bayesian approach, where subjective (judgmental, knowledge-based) probabilities are used to express the epistemic uncertainties about the parameters. This type of uncertainty analysis is an integrated part of risk assessment.

However, uncertainty analysis also exists independently of risk assessment (Morgan and Henrion, 1990). Formally, uncertainty analysis refers to the determination of the uncertainty associated with the results of an analysis that derives from uncertainty related to the input to the analysis (including the methods and models used in the analysis) (Helton *et al.*, 2006).

We may illustrate the ideas of uncertainty analysis by introducing a model $g(X)$, which depends on the input quantities X and on the function g. The quantity of interest, Z, is computed by using the model $g(X)$. The uncertainty analysis of Z requires an assessment of the uncertainties about X and their propagation through the model g to produce an assessment of the uncertainties concerning Z, see Figure 1.6. Uncertainty related to the model structure g, that is, uncertainty about the error $Z - g(X)$, is typically treated separately (Devooght, 1998; Zio and Apostolakis, 1996; Baraldi and Zio, 2010). In fact, while the impact of uncertainties associated with X has been widely investigated and many sophisticated methods have been developed to deal with it, research is still ongoing to obtain effective and agreed methods to handle the uncertainty related to the model structure (Parry and Drouin, 2009). For a broad sample of methods and ideas concerning model uncertainty, see Mosleh *et al.* (1994).

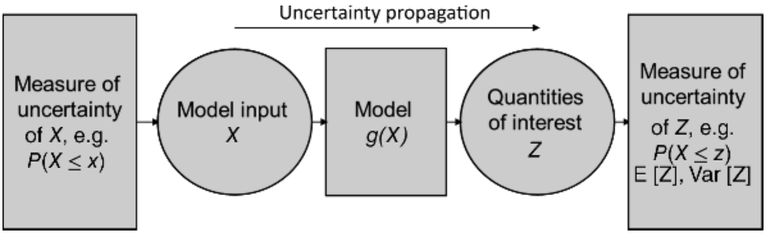

Figure 1.6 Basic features of uncertainty analysis (based on de Rocquigny, Devictor, and Tarantola (2008), Helton et al. *(2006), and Aven (2010a)).*

Overall frameworks for uncertainty analysis and management have been developed based on the elements Z and $g(X)$ as indicated above; see, for example, de Rocquigny, Devictor, and Tarantola (2008), Helton *et al.* (2006), and Aven (2010a).

These frameworks also applies to risk assessment. Then Z could for example be some high-level event of interest, such as a blowout in an offshore QRA setting or a core meltdown in a nuclear PRA setting, and X could be the set of low-level events which through various combinations could lead to the occurrence of the high-level event.

The quantities X and Z could also be frequentist probabilities representing fractions in a large (in theory, infinite) population of similar items, that is, parameters of probability models. Think of the frequentist probability p introduced in Section 1.1.3 that the person gets a specific disease. In this case, the assessment is consistent with the probability of frequency approach briefly outlined at the end of Section 1.1.3, which is based on the use of judgmental probabilities to express epistemic uncertainty about unknown frequentist probabilities (Kaplan and Garrick, 1981).

Finally, we add a note on sensitivity analysis, which is not the same as uncertainty analysis although they are closely linked. Sensitivity analysis indicates how sensitive the considered metrics are with respect to changes in basic input quantities (e.g., parameter values, assumptions and suppositions made) (Saltelli *et al.*, 2008; Cacuci and Ionescu-Bujor, 2004; Frey and Patil, 2002; Helton *et al.*, 2006). In an uncertainty analysis context, more specific definitions of sensitivity analysis have been suggested: for example, according to Helton *et al.* (2006), sensitivity analysis refers to the determination of the contributions of individual uncertain inputs to the analysis results.

In engineering risk assessments, a distinction is commonly made between aleatory and epistemic uncertainty (e.g., Apostolakis, 1990; Helton and Burmaster, 1996) as mentioned above in relation to the Bayesian approach, see Figure 1.7. Aleatory uncertainty refers to variation in populations, and epistemic uncertainty to lack of knowledge about phenomena. The latter usually translates into uncertainty about the parameters of a probability model used to describe variation. Whereas epistemic uncertainty can be reduced, aleatory uncertainty cannot, and it is sometimes called irreducible uncertainty (Helton and Burmaster, 1996). The aleatory uncertainty concept

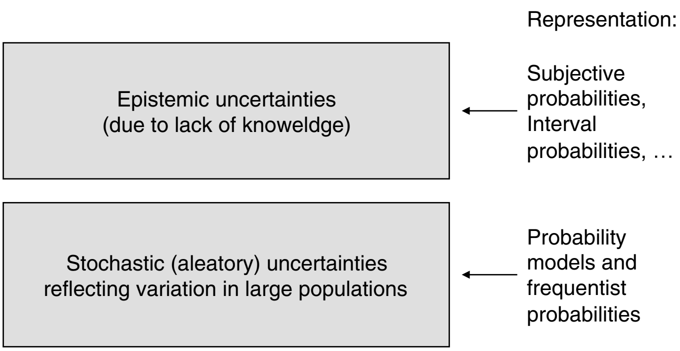

Figure 1.7 Illustration of the two types of uncertainties, aleatory and epistemic, and how they are represented.

is variably referred to as stochastic uncertainty (e.g., Helton and Burmaster, 1996), randomness, (random) variation (e.g., Aven, 2012a) or (random) variability (e.g., Baudrit, Dubois, and Guyonnet, 2006). It requires a large (in theory, infinite) population of "similar" but not identical units (realizations) to become operational, for example, a population of mass-produced units, a population of human beings or a sequence of dice throws. The epistemic uncertainty concept is variously referred to as subjective uncertainty (e.g., Helton and Burmaster, 1996), (partial) lack of knowledge (e.g., Aven, 2012a), or partial ignorance (e.g., Ferson and Ginzburg, 1996).

Returning to the health risk example of Section 1.1.3, the binomial probability model and the frequentist probability p represent the aleatory uncertainty, the variation in populations, whereas the subjective probability distribution $H(p')$ describes the epistemic uncertainties, reflecting the assessor's judgment about the true value of p.

1.5 Challenges: Discussion

In all the examples above we have used probabilities to describe risk. However, there are a growing number of researchers and analysts who find the probability-based approaches for assessing risk and uncertainties to be too narrow (see, e.g., Renn, 1998; Aven and Zio, 2011; Aven, 2011b). The argumentation follows different lines of thinking. One main point raised is that risk is more than some analysts' subjective probabilities, which may be based on strong assumptions and lead to poor predictions of Z. One may assign a low probability of health problems caused by the use of some new chemicals, but these probabilities could produce poor predictions of the actual number of people that experience such problems. Or one may assign a probability of fatalities occurring on an offshore installation based on the assumption that the installation structure will withstand a certain accidental collision energy load while, in reality, the structure could fail at a lower load level: the assigned probability did not reflect this uncertainty.

While probabilities can always be assigned under the subjective probability approach, the origin and amount of information supporting the assignments are not reflected by the numbers produced. One may, for example, subjectively assess that two different events have probabilities equal to, say, 0.7, but in one case the assignment is supported by a substantial amount of relevant data, whereas in the other by effectively no data at all. This is the main argument in the critique of the probability-based approach to risk and uncertainty.

Another important argument relates to the decision setting. In a risk assessment context there are often many stakeholders and they may not be satisfied with a probability-based assessment expressing the subjective judgments of the analysis group; a broader risk description is sought.

Probability models constitute a pillar of the probabilistic approach, an essential tool for assessing uncertainties and drawing useful insights (Helton, 1994; Winkler, 1996). The probability models coherently and mechanically facilitate the updating of probabilities. However, for many types of applications these models cannot be justified, the consequence being that the probability-based approach to risk and uncertainty becomes difficult to implement. A probability model presumes some sort of model stability; populations of similar units need to be constructed. But such stability is often not fulfilled (Bergman, 2009; Aven and Kvaløy, 2002). Consider the definition of a chance. In the case of a die, we would establish a probability model expressing that the distribution of outcomes is given by $(p_1, p_2, \ldots p_6)$, where p_i is the chance of outcome i, interpreted as the fraction of throws resulting in outcome i. However, in a risk assessment context the situations are often unique and the establishment of chances means the construction of fictional populations of non-existing similar situations. Then, chances and probability models in general cannot be easily defined as in the die-tossing example; in many cases they cannot be meaningfully defined at all. For example, it makes no sense to define a chance (frequentist probability) of a terrorist attack (Aven and Renn, 2010, p. 80). In other cases the conclusion may not be so obvious. For example, the chance of an explosion in a process plant may be introduced in a risk assessment, although the underlying population of infinite similar situations is somewhat difficult to describe.

1.5.1 Examples

1.5.1.1 LNG plant risk

An LNG (Liquefied Natural Gas) plant is planned and the operator would like to locate the plant no more than a few hundred meters from a residential area (Vinnem, 2010). Risk assessments are performed demonstrating that the risk is acceptable according to some predefined risk acceptance criteria. Risk is expressed using computed probabilities and expected values. Both IR numbers and *f–n* curves are presented. However, the assessments and the associated risk management are subject to strong criticism. The neighbors as well as many independent experts do not find the risk characterization sufficiently informative to support the decision making on the location and design of the plant. Sensitivity analyses are lacking, as well as reflections on uncertainties. The risk figures produced are based on a number of critical assumptions, but these assumptions are neither integrated into the risk characterization presented, nor communicated by the operator.

This type of risk assessment has been conducted for many years (Kolluru *et al.*, 1996) but the problems seem to be the same. Still, a narrow risk characterization is provided. One reason may be the fact that acknowledging uncertainty can weaken the authority of the decision maker and agency, by creating an image of being unknowledgeable (Tickner and Kriebel, 2006). We will not, however, discuss this any further here; the point we would like to make is simply that these problems are also a result of the perspective adopted for the risk assessments. The conventional view, supported by most authoritative guides on risk assessments (e.g., Bedford and Cooke, 2001; Vose, 2008), is to a large extent pleased with the "narrow" probabilistic-based assessments as in the LNG case. Based on this view, the assessments may be criticized for too few and too restricted sensitivity and uncertainty analyses, but not for systematic hiding or camouflaging uncertainties.

1.5.1.2 Terrorism risk

For many of the currently emerging risks, the need for an extended approach is even more urgent. Take for example terrorism risk. Here the risk assessments often focus on estimating the probability of an attack, P(attack). But what is the meaning of such a probability? The conventional assessment is based on probability models, so there is a need for the definition of a frequentist probability or chance $p = P$(attack). But such interpretations are meaningless, as mentioned above (Aven and Renn, 2010, p. 80). To define such a probability (chance) we need to construct an infinite population of similar attack situations. The proportion of successes equals the probability of an attack. But it makes no sense to define a large set of "identical," independent attack situations, where some aspects (e.g., related to the potential attackers and the political context) are fixed and others (e.g., the attackers' motivation) are subject to variation. Say that the attack success rate is 10%. Then, in 1000 situations, with the attackers and the political context specified, the attackers will successfully attack in about 100 cases. In these situations the attackers are motivated, but not in the remaining ones. Motivation for an attack in one situation does not affect the motivation in another. For independent repeated random situations such "experiments" are meaningful, but not in unique cases like this. Still, many researchers and analysts work within such a risk assessment framework.

1.5.2 Alternatives to the probability-based approaches to risk and uncertainty assessment

Based on the above critiques, it is not surprising that alternative approaches for representing and describing risk and uncertainties have been suggested; see, for example, the special issues on imprecision in the *Journal of Statistical Theory and Practice* (JSTP, 2009), on uncertainty in engineering risk and reliability in the *Journal of Risk and Reliability* (JRR, 2010), and on alternative representations of epistemic uncertainty in *Reliability Engineering and System Safety* (RESS, 2004). Four main categories are (Dubois, 2010; Aven and Zio, 2011):

a. Probability-bound analysis, combining probability analysis and interval analysis (Ferson and Ginzburg, 1996). Interval analysis is used for those components whose aleatory uncertainties cannot be accurately estimated; for the other components, traditional probabilistic analysis is carried out.

b. Imprecise probability, after Walley (1991), and the robust statistics area (Berger, 1994) (see also Coolen, Augustin, and Troffaes, 2010; Klir, 2004), which encompasses probability-bound analysis, and certain aspects of evidence and possibility theory as special cases.

c. Evidence theory (or belief function theory), as proposed by Dempster (1967) and Shafer (1976), and the closely linked theory of random sets (Nguyen, 2006).

d. Possibility theory (Dubois and Prade, 1988, 2009; Dubois, 2006), which is formally a special case of the imprecise probability and random set theories.

We will thoroughly review and discuss these approaches in this book, as well as attempts made to combine different approaches, for example, probabilistic analysis and possibility theory, where the uncertainties of some parameters are represented by probability distributions and the uncertainties of the remaining parameters by means of possibility distributions; see Baudrit, Dubois, and Guyonnet (2006) and the applications in Baraldi and Zio (2008), Flage *et al.* (2013), and Helton, Johnson, and Oberkampf (2004).

All these approaches and methods produce epistemic-based uncertainty descriptions and in particular intervals, but they have not been broadly accepted by the risk assessment community. Much effort has been made in this area, but there are still many open questions related to the foundation of these approaches and their use in risk and uncertainty decision making; see the discussions in, for example, Cooke (2004), Smets (1994), Lindley (2000), Aven (2011a), Bernardo and Smith (1994), and North (2010). Many risk researchers and risk analysts are skeptical about the use of the alternative approaches (such as those of the four categories (a)–(d) mentioned above) for the representation and treatment of uncertainty in risk assessment for decision making, and some also argue intensively against them; see, for example, North (2010, p. 380).

However, as argued above, the probability-based approach does not solve the problems raised. The decision basis cannot be restricted to assigned probabilities: there is a need to go beyond the traditional Bayesian approach; a broader perspective and framework is required. The present book is based on such a conviction. As a matter of fact, no comprehensive authoritative guidance exists today on when to use probability and when to use an alternative representation of uncertainty. The main challenge is to define the conditions when probability is the appropriate representation of uncertainty. An argument often propounded is that probability is the appropriate representation of uncertainty only when a sufficient amount of data exists on which to base the probability (distribution) in question. But it is not obvious how to make such a prescription operational (Flage, 2010). Consider the representation of uncertainty about the parameter(s) of a probability model. If a large enough amount of data exists, there would be no uncertainty about the parameter(s) and hence no need for a representation of such uncertainty. When is there enough data to justify probability, but not enough to accurately specify the true value of the parameter in question and, thus, make probability, as an epistemic concept, superfluous?

Other approaches for representing the uncertainties have also been suggested. One example is the approach based on the maximum entropy principle. This approach does not require the specification of the whole probability distribution but only of some of its features, for example, the mean and variance; then, a mathematical procedure is applied to obtain the distribution characterized by the specified features and, in a certain sense, minimum information beyond that, see Bedford and Cooke (2001). Another approach relevant in this context is probabilistic inference with uncertain and partial evidence, developed by Groen and Mosleh (2005) as a generalization of Bayes' theorem.

Taking a broader view, we may identify different directions of development with respect to (alternative) representations of uncertainty in risk analysis (Flage, 2010).

One direction, as suggested by Lindley (2006) and O'Hagan and Oakley (2004), is to retain probability as the sole representation of uncertainty and to focus on improving the measurement procedures for probability. Another direction is that resulting in a semi-quantitative approach, where quantitative risk metrics are supported by qualitative assessments of the strength of background knowledge of these metrics (Aven, 2013). Such an approach presupposes an acknowledgment and belief that the full scope of risk and uncertainty cannot be transformed into a mathematical formula, using probability or any other measure of uncertainty. Numbers can be generated, but these alone would not serve the purpose of the risk assessment, to reveal and describe the risks and uncertainties. Furthermore, a duality in terms of interpretation, like that which affects probability (limiting relative frequency vs. degree of belief), also affects possibility theory (degree of compatibility or ease vs. lower and upper probability) and the theory of belief functions (degree of belief per se vs. lower and upper probability) (Flage, 2010). Another direction thus consists of assessing the appropriateness of lower and upper probability vs. other interpretations for risk analysis purposes, and then developing a proper foundational basis like those mentioned above in the Bayesian setting. Finally, there is the "unifying" approach, based on the combination of different representations. These development directions are further studied in Chapter 7.

1.5.3 The way ahead

When considering the methods for representing and characterizing uncertainties in risk assessment, two main concerns need to be balanced:

1. Knowledge should, as far as possible, be "inter-subjective" in the sense that the representation corresponds to "documented and approved" information and knowledge ("evidence"); the methods and models used to treat this knowledge should not add information that is not there, nor ignore information that is there.
2. Analysts' judgments ("degrees of belief") should be clearly reflected ("judgments").

The former concern can make the pure Bayesian approach difficult to apply in certain instances: when scarce information and little knowledge are available, introducing analysts' subjective probability distributions may be unjustifiable since this leads to building a structure in the analysis that is not present in the information and knowledge. For example, if an expert states his or her uncertainty assessment on a parameter value in terms of a range of possible values, this does not justify the allocation of a specific probability distribution function (e.g., the uniform probability distribution) onto the range. In this view, it might be said that a more defense-in-depth (bounding) representation of the information and knowledge available would be one which leaves the analysis open to all possible probability distribution structures on the assessed range, without imposing one in particular and without excluding any, thus providing results which bound all possible distributions.

At the same time, the representation framework should also take into account the second concern above, that is, allow for the transparent inclusion of preferential assignments by the experts (analysts) who wish to express that, according to their beliefs, some values are more or less likely than others. The Bayesian approach is the proper framework for such assignments.

From the point of view of the quantitative modeling of uncertainty in risk assessment, two topical issues are the proper handling of dependencies among uncertain parameters, and of model uncertainties. No matter what modeling paradigm is adopted, it is critical that the meaning of the various concepts be clarified. Without such clarification it is impossible to build a scientifically based risk assessment. In complex situations, when the propagation is based on many parameters, strong assumptions may be required to be able to carry out the analysis in practice. The risk analysts may acknowledge a degree of dependency, but the analysis may not be able to describe it in an adequate way. The derived uncertainty representations must be understood and communicated as measures conditional on this constraint. In practice it is the main task of the analysts to seek simple representations of the system performance, and by smart modeling it is often possible to obtain independence. The models used are also included in the background knowledge of epistemic-based uncertainty representations. We seek accurate models, but at the same time simple models. The choice of the right model cannot be seen in isolation from the purpose of the risk assessments.

Acknowledging risk and uncertainty assessments as decision support tools requires that the meaning and practical interpretation of the quantities computed are presented and communicated in an understandable format to the decision makers. The format must allow for meaningful comparisons to numerical safety criteria, if defined, for manipulation (e.g., by screening, bounding, and/or sensitivity analyses) and for communication in deliberation processes. This issue has been addressed by many researchers in the scientific literature; see, for example, the recent discussions in Aven (2010b), Dubois (2010), and Dubois and Guyonnet (2011) as well as Renn (2008). However, there are still many questions that remain to be answered, for example, concerning the type of information the decision maker needs in order to be risk informed, as the debate between Aven (2010b) and Dubois (2010) reveals. We will address this issue in detail in the coming chapters.

In the coming chapters, we will perform a detailed review and discussion of the most common and relevant approaches and methods for representing and characterizing uncertainties in risk assessment. The review and discussion are based on the perspective and framework of risk assessment and uncertainty analysis introduced in this chapter. This perspective and framework extend beyond the Bayesian approach. We have argued that a full risk–uncertainty description is more than subjective probabilities. Risk is about hazards/threats, their consequences, and the associated uncertainties, and to assess risk various tools can be used to measure the uncertainties. In the coming chapters, we will look closer at the most important of these tools, to provide the reader with an improved basis for selecting appropriate approaches and methods for representing and characterizing uncertainties in risk assessment.