1

Introduction to risk management and risk assessments. Challenges

This chapter provides a broad introduction to risk management and risk asssessment, as a basis for the analyses and dicussions in the coming chapters. The presentation highlights general features but also challenges related to the definitions and use of these tools. Key references for the chapters are Bedford and Cooke (2001), Vose (2008) and Aven and Vinnem (2007). The terminology is to a large extent in line with ISO (2009a). See summary of key definitions in Appendix B.

1.1 General features of risk management and risk assessments

Risk management is all coordinated activities to direct and control an organisation with regard to risk. Two main purposes of the risk management are to ensure that adequate measures are taken to protect people, the environment and assets from undesirable consequences of the activities being undertaken, and to balance different concerns, for example safety and costs. Risk management covers both measures to avoid the occurrence of hazards/threats and measures to reduce their potential consequences. In industries like nuclear and oil & gas, risk management was traditionally based on a prescriptive regulating regime, in which detailed requirements for the design and operation of the plant were specified (Kumamoto, 2007; Aven and Vinnem, 2007). This regime has gradually been replaced by more goal-oriented regimes, putting emphasis on what to achieve rather than on the means of doing so. Goal orientation and risk characterisations are two major components of these new regimes that have been enthusiastically endorsed by international organisations and various industries (see e.g. IAEA Guidelines (1995), HSE (2001), Kröger (2006); the IPCS and WHO risk terminology document (2004) and the risk management guidelines of the EU Commission (European Commission, 2000, 2003; IEC, 1993)). Such

an approach to risk management is believed to provide higher levels of performance both in terms of productivity and risk reduction (Aven and Renn, 2009b).

Quantitative risk assessment

Quantitative Risk Assessment (QRA) (also referred to as Probabilistic Risk Assessment – PRA) is a key tool used in these new approaches. A QRA systemises the present state of knowledge including the uncertainties about the phenomena, processes, activities and systems being analysed. It identifies possible hazards/threats (such as a gas leakage or a fire), analyses their causes and consequences, and describes risk. A QRA provides a basis for characterising the likely impacts of the activity studied, for evaluating whether risk is tolerable or acceptable and for choosing the most effective and efficient risk policy, for example with respect to risk-reducing measures. It allows for the calculation of expected values so that different risks can be directly compared. Common practice in probabilistic risk assessment avoids, however, the aggregation of the two components and leaves it to the risk evaluation or management team to draw the necessary conclusions from the juxtaposition of loss and probabilities (Aven, 2003; Kröger, 2005). In addition, second-order uncertainties are introduced via different types of uncertainty intervals to make the confidence of probability judgements more explicit (Apostolakis and Pickett, 1998; Aven, 2003), see also Sections 2.7 and 8.3. For some extensive reviews of the use of QRA/PRA in a historical perspective, see Rechard (1999, 2000).

Some of the basic tools used for analysing the probabilities and risk are statistical estimation theory, fault tree analysis (FTA) and event tree analysis (ETA). These tools belong to the following main categories of basic analysis methods:

(a) *Statistical methods:* Data are available to predict the future performance of the activity or system analysed. These methods can be based on data extrapolation or probabilistic modelling.
(b) *Systems analysis methods:* These methods (which include FTA and ETA) are used to analyse systems where there is a lack of data to accurately predict the future performance of the system. Insights are obtained by decomposing the system into subsystems/components for which more information is available. Overall probabilities and risk are a function of the system's architecture and of the probabilities on the subsystems/component level (Paté-Cornell and Dillon, 2001).

Quantitative risk assessment (QRA) is often associated with system analysis methods (see e.g. Bedford and Cooke, 2001), but in this book we interpret QRA (PRA) as any risk assessment which is based on quantification of risk using probabilities.

A number of new and improved methods have been developed in recent years to better meet the needs of the analysis, in light of the increasing complexity of the systems and to respond to the introduction of new technological systems (Aven and Zio, 2011). Many of the methods introduced allow for increased levels of detail and precision in the modelling of phenomena and processes within an integrated framework of analysis covering physical phenomena, human and organisational factors as well as software dynamics (e.g. Mohaghegh *et al*., 2009; Luxhoj *et al*., 2001; Ale *et al*., 2009; Røed *et al*., 2009). Other methods are devoted to the improved representation and analysis of the risk and related uncertainties, in view of the decision-making tasks that the outcomes of the analysis are intended to support. Examples of relatively newly introduced methods are Bayesian Belief Networks (BBNs), Binary Digit Diagrams (BDDs), multi-state reliability analysis, Petri Nets 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 traditional risk assessment approach used in QRAs can be viewed as a special case of system engineering (Haimes, 2004). This approach, which to a large extent is based on causal chains and event modelling, has been subject to strong criticism (e.g. Rasmussen, 1997; Hollnagel, 2004; Leveson, 2004). It is argued that some of the key methods used in risk assessments are not able to capture "systemic accidents". Hollnagel (2004), for example, argues that to model systemic accidents it is necessary to go beyond the causal chains – we must describe system performance as a whole, where the steps and stages on the way to an accident are seen as parts of a whole rather than as distinct events. It is not only interesting to model the events that lead to the occurrence of an accident, which is done for example in event and fault trees, but also to capture the array of factors at different system levels that contribute to the occurrence of these events. Leveson (2007) makes her points very clear:

> Traditional methods and tools for risk analysis and management have not been terribly successful in the new types of high-tech systems with distributed human and automated decision-making we are attempting to build today. The traditional approaches, mostly based on viewing causality in terms of chains of events with relatively simple cause-effect links, are based on assumptions that do not fit these new types of systems: These approaches to safety engineering were created in the world of primarily mechanical systems and then adapted for electro-mechanical

systems, none of which begin to approach the level of complexity, non-linear dynamic interactions, and technological innovation in today's socio-technical systems. At the same time, today's complex engineered systems have become increasingly essential to our lives. In addition to traditional infrastructures (such as water, electrical, and ground transportation systems), there are increasingly complex communication systems, information systems, air transportation systems, new product/process development systems, production systems, distribution systems, and others.

Leveson (2004) argues for a paradigm-changing approach to safety engineering and risk management. She refers to a new alternative accident model, called STAMP (System-Theoretic Accident Modeling and Processes).

Nonetheless, the causal chains and event modelling approach has shown to work for a number of industries and settings. It is not difficult to point at limitations of this approach, but the suitability of a model always has to be judged with reference to not only its ability to represent the real world, but also its ability to simplify the world. All models are wrong, but they can still be useful to use a well-known phrase. Furthermore, the causal chains and event modelling approach is continuously improved, incorporating human, operational and organisational factors, as was mentioned above. Mohaghegh *et al*. (2009), for example, present a "hybrid" approach for analysing dynamic effects of organisational factors on risk for complex socio-technical systems. The approach links system dynamics, Bayesian belief networks, event sequence diagrams and fault trees.

For the purpose of the present book, it suffices to consider the basic analysis tools such as fault tree and event tree models, probability models and statistical inference based on these models.

Risk assessment covers risk analysis and risk evaluation; see Figure 1.1. Risk analysis is a methodology designed to determine the nature and extent of risk. It comprises the following three main steps:

1. Identification of hazards/threats/opportunities (sources)
2. Cause and consequence analysis, including analysis of vulnerabilities
3. Risk description, using probabilities and expected values.

This definition of risk analysis seems to be the most common, but there are others (refer to IRGC, 2005). One of these considers risk analysis as an overall concept, comprising risk assessment, risk perception, risk management, risk communication, and their interactions. This interpretation has been often used among members of the Society of Risk Analysis.

Expressing risk also means to perform sensitivity analyses. The purpose of these analyses is to show how sensitive the output risk indices are with respect to changes in basic input quantities, assumptions and suppositions.

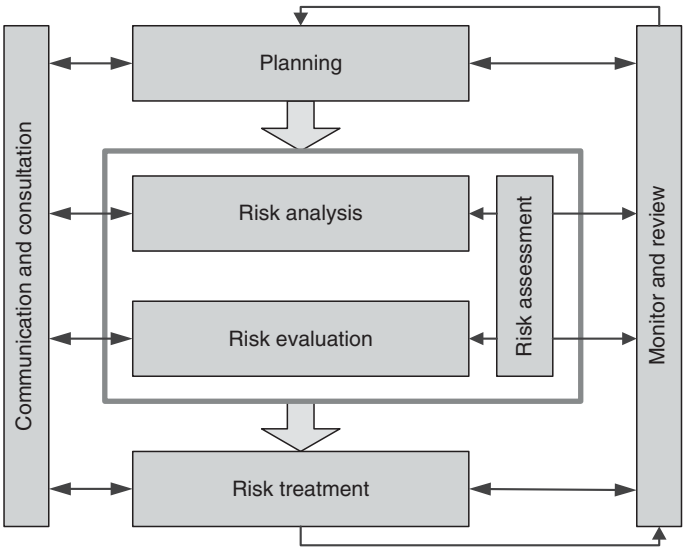

Figure 1.1 The risk assessment process (based on ISO, 2009b). Note that the ISO (2009a,b) does not include source identification as a part of risk analysis.

The sensitivity analyses can be used to identify critical systems, and thus provide a basis for selecting appropriate measures. To illustrate this, let R be a risk index, for example expressing the expected number of fatalities or the probability of a system failure, and let R_i be the risk index when subsystem i is in the functioning state. Then a common way of ranking the different subsystems is to compute the risk improvement potential (also referred to as the risk achievement worth) $I_i = R_i - R$, i.e. the maximum potential risk improvement that can be obtained by improving system i. The potential I_i is referred to as a risk importance measure. See Aven and Nøkland (2010) for a recent review of such measures.

Having established a risk description (risk picture), its significance is then evaluated (risk evaluation). Is the risk high compared to relevant reference values or decision criteria? How does alternative A compare with alternative B? etc. Risk analysis is often used in combination with risk acceptance criteria, as inputs to risk evaluation. Sometimes the term "risk tolerability limits" is used instead of risk acceptance criteria. The criteria state what is deemed as an unacceptable risk level. The need for risk-reducing measures is assessed with reference to these criteria. In some industries and countries, it is a requirement in regulations that such criteria should be defined in advance of performing the analyses.

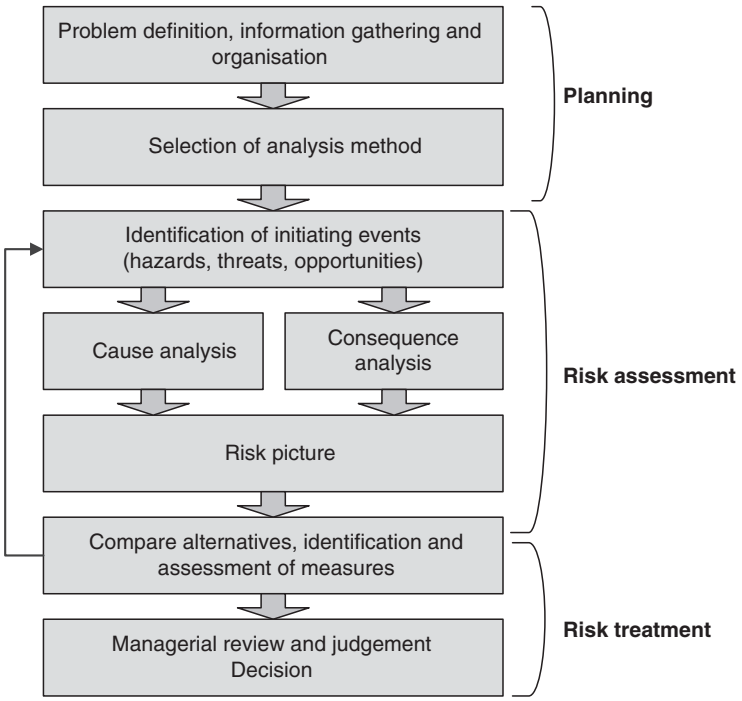

Figure 1.2 The main steps of the risk assessment process, covering the planning, the risk assessment and its use (based on Aven, 2008a).

The risk assessment process (planning, execution and use of risk assessments)

Risk assessment is followed by risk treatment, which is a process involving the development and implementation of measures to modify risk, including measures designed to avoid, reduce ("optimise"), transfer or retain risk. Risk transfer means sharing with another party the benefit or loss associated with the risk. It is typically effected through insurance.

"Planning" defines the basic frame conditions within which the risks must be managed and sets the scope for the rest of the risk assessment process. It means definition of suitable decision criteria as well as structures for how to carry out the risk assessment.

It is possible to detail the process in Figure 1.1 in many different ways to illustrate the planning, execution and use of risk analyses. Figure 1.2 shows an example based on Aven (2008a).

The results of the assessments need to be evaluated in the light of the premises, assumptions and limitations of these assessments. We refer to this stage of the process as the managerial review and judgement (Hertz and

Thomas, 1983; Aven, 2003). The assessments are based on some background knowledge that must be reviewed together with the results of the assessments. Consideration should be given to factors such as (Aven, 2003):

- which decision alternatives have been analysed
- which performance measures have been assessed
- the fact that the results of the analyses represent judgements (expert judgements)
- difficulties in assigning probabilities in the case of large uncertainties
- the fact that the assessments' results apply to models that are simplifications of the real world and real world phenomena.

The decision-making basis will seldom be in a format that provides all the answers that are important to the decision-maker. There will always be limitations in the information basis and the review and judgement described means that one views the basis in a larger context. Perhaps the analysis did not take into consideration what the various measures mean for the reputation of the enterprise, but this is obviously a condition that is of critical importance for the enterprise. The review and judgement must also cover this aspect.

The weight the decision-maker gives to the basis information provided depends on the confidence he/she has in those who developed this information. However, even if the decision-maker has maximum confidence in those doing this work, the decision still does not come about on its own. It is often difficult to make decisions when the risk is high. The decisions encompass difficult considerations and weighting with respect to uncertainties and values, and this cannot be delegated to those who create the basis information. It is the responsibility of the decision-maker to undertake such considerations and weighting, and to make a decision that balances the various concerns.

Apostolakis (2004, p. 518) makes this clear:

> I wish to make one thing very clear: QRA results are *never* the sole basis for decision-making by responsible groups. In other words, safety-related decision-making is *risk-informed*, not risk-based.

Figure 1.3 illustrates the use of risk assessment in the decision-making. Risk assessment is carried out to support the decision-making, for example a choice between various concepts, design configurations, risk-reducing measures etc. Other types of assessment are also needed, such as cost-effectiveness analyses and cost–benefit analyses.

The same types of ideas are reflected in many other decision analysis frameworks and contexts, for example the analytic-deliberative process

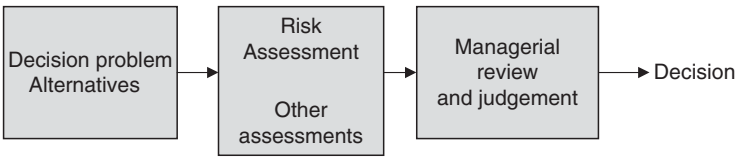

Figure 1.3 Model of the use of risk assessment to support decision-making.

recommended by the US National Research Council (1996) in environmental restoration decisions involving multiple stakeholders. According to this process, analysis "uses rigorous, replicable methods, evaluated under the agreed protocols of an expert community – such as those of disciplines in the natural, social, or decision sciences, as well as mathematics, logic, and law – to arrive at answers to factual questions"; while "deliberation is any formal or informal process for communication and collective consideration of issues. ... Participants in deliberation discuss, ponder, exchange observations and views, reflect upon information and judgements concerning matters of mutual interest and attempt to persuade each other." Such a process is particularly adapted to and relevant to decisions of great public interest.

Various decision-making strategies can form the basis for the decision. By "decision-making strategy" we mean the underlying thinking that goes on, and the principles that are to be followed with respect to how the decision is to be made, and how the process prior to the decision should be. Central to this is the question of who will be involved, how to use the various forms of analyses, and how the actual process is to be carried out.

ALARP principle

An example of such a strategy is to use risk acceptance (tolerability) criteria as inputs to risk evaluation. Another strategy is to adopt the ALARP principle, which means that risk should be reduced to a level that is as low as reasonably practicable. According to the ALARP principle, a risk-reducing measure should be implemented provided it cannot be demonstrated that the costs are grossly disproportionate relative to the gains obtained (the burden of proof is reversed). The standard approach when applying the ALARP principle, as for example used in the UK, is to consider three regions:

1. the risk is so low that it is considered negligible
2. the risk is so high that it is intolerable
3. an intermediate level where the ALARP principle applies.

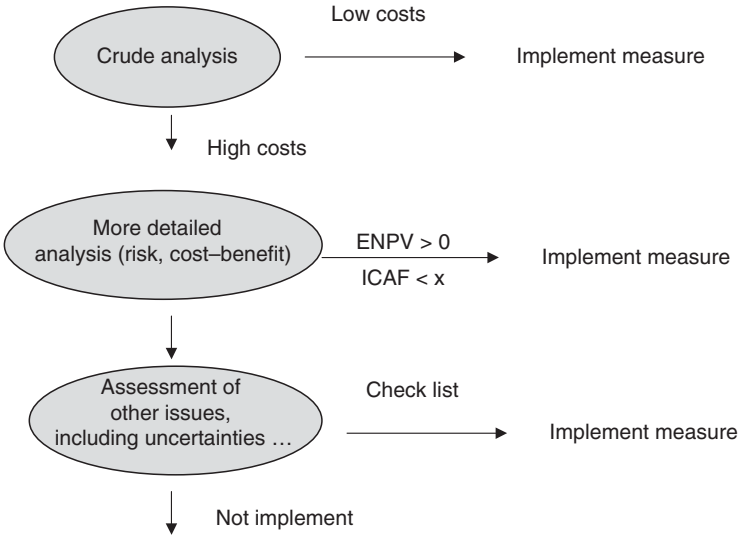

Figure 1.4 Procedure for implementing ALARP and the gross disproportionate criterion (Aven and Vinnem, 2007).

In most cases in practice risk is found to be in region 3 and the ALARP principle is adopted. This will include a dedicated search for possible risk-reducing measures and a subsequent assessment of these in order to determine which to be implemented.

To verify ALARP, procedures mainly based on engineering judgements and codes are used, but also traditional cost–benefit analyses and cost-effectiveness analyses. When using such analyses, guidance values as above are often used to specify what values define "gross disproportion".

Conclusions are often self-evident when computing indices such as the expected cost per expected number of lives saved. For example, a strategy may be that measures will be implemented if the expected cost per expected number of lives saved (Implied Cost of Averting a Fatality – ICAF) is less than €2 million. Figure 1.4 sketches the main ideas of a procedure for how to implement ALARP and the gross disproportionate criterion in practice presented in Aven and Vinnem (2007).

The procedure can be summarised as follows:

- Perform a crude qualitative analysis of the benefits and burdens of the risk-reducing measure. If the costs are not judged to be large, implement the measure. Gross disproportion has not been demonstrated.
- If the costs are considered large, quantify the risk reduction and perform an economic analysis as indicated above (computing for example ICAF or the

expected net present value, i.e. E[NPV]). If E[NPV] > 0 or ICAF is low (typically less than some few € millions), implement the measure. Gross disproportion has not been demonstrated.

If these criteria are not met, assess uncertainty factors and other issues of relevance not covered by the previous analyses. A checklist is used for this purpose. Aspects that could be covered by this list are:

- Is there considerable uncertainty (related to phenomena, consequences, conditions) and will the measure reduce these uncertainties?
- Does the measure significantly increase manageability? High competence among the personnel can give increased assurance that satisfactory outcomes will be reached.
- Is the measure contributing to obtaining a more robust solution?
- Is the measure based on best available technology (BAT)?
- Are there unsolved problem areas: personnel safety-related and/or work environment-related?
- Are there possible areas where there is conflict between these two aspects?
- Is there a need for strategic considerations?

If the risk-reducing measure scores high on these factors (many yes answers), gross disproportion has not been demonstrated.

- Otherwise, the costs are in gross disproportion to the benefits gained, and the measures should not be implemented.

Cautionary and precautionary principles

The ALARP principle can be considered as a special case of the cautionary principle which states that in the face of uncertainty and risk, *caution* should be a ruling principle, for example by not starting an activity, or by implementing measures to reduce risks and uncertainties (HSE, 2001; Aven and Vinnem, 2007, p. 34). This principle is being implemented in all industries through safety regulations and requirements. For example, in the Norwegian petroleum industry it is a regulatory requirement that the living quarters on an installation should be protected by fireproof panels of a certain quality, for walls facing process and drilling areas. This is a standard adopted to obtain a minimum safety level. It is based on established practice of many years of operation of process plants. A fire may occur; it represents a hazard for the personnel and in the case of such an event, the personnel in the living quarters should be protected. The assigned probability for the living quarter on a specific installation being exposed to fire may be judged as low, but we know