
RISK ANALYSIS AND MANAGEMENT. BASIC CONCEPTS AND PRINCIPLES

T. Aven

University of Stavanger, Norway

e-mail: terje.aven@uis.no

ABSTRACT

This paper reviews and discusses some key concept and principles of risk analysis and risk management, based on a set of statements, formulated as myths about risk. Examples of such myths are: risk is equal to the expected value, risk equals uncertainty, risk can be expressed by probabilities, risk is equal to an event, risk acceptance criteria contribute to obtaining a high safety standard, and ALARP can be verified by cost-benefit analyses. It is concluded that risk needs to address both the consequences and the uncertainties about the consequences, and that it is necessary to see beyond expected values and probabilities.

1 INTRODUCTION

There is an enormous drive and enthusiasm in various industries, services and society as a whole nowadays to implement risk management in the organizations. There are high expectations, that risk management is the proper framework for obtaining high levels of performance. We see a lot of initiatives to establish adequate concepts and tools. However, the risk management discipline is young, and there are many difficult issues and challenges. These relate in particular to the foundation and use of risk analyses; how to express risk, how to handle uncertainties, and how to use risk analysis in a decision-making context. These issues are addressed in this paper. The purpose of the paper is to review and discuss some key concept and principles of risk analysis and risk management. We do this by formulating a set of statements, which can be seen as *myths* about risk. These myths are presented and discussed in the following section. Some conclusions are provided in Section 3. Risk management is defined as all measures and activities carried out to manage risk. Risk management deals with balancing the conflicts inherent in exploring opportunities on the one hand, and avoiding losses, accidents, and disasters, on the other (Aven & Vinnem 2007).

To support decision-making on design and operation, risk analyses are conducted. The analyses include identification of hazards and threats, cause analyses, consequence analyses and risk description. The results of the analyses are then evaluated. The totality of the analyses and the evaluations are referred to as risk assessment. 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 (“optimize”), transfer or retain risk. Risk transfer means sharing with another party the benefit or loss associated with a risk. It is typically affected through insurance. The terminology is in line with the ISO standard on risk management terminology (ISO 2002).

By carrying out a risk analysis one can:

- Establish a risk picture
- Compare different alternatives and solutions in terms of risk
- Identify factors, conditions, activities, systems, components, etc. that are important (critical) with respect to risk
- Demonstrate the effect of various measures on risk.

This provides a basis for:

- Choosing between various alternative solutions and activities while in the planning phase of a system
- Choosing between alternative designs of a solution or a measure.
- Drawing conclusions on whether specific solutions and measures meet stated requirements
- Setting requirements for various solutions and measures, for example related to the performance of the preparedness systems
- Documenting an acceptable safety and risk level.

2 MYTHS ABOUT RISK

We will discuss the following myths about risk:

1. Risk is equal to the expected value
2. Risk equals uncertainty
3. Risk is equal to an event
4. Risk is a probability or a probability distribution
5. Risk equals expected disutility
6. Risk is restricted to the case of known probabilities
7. Risk based on subjective probabilities is the same as risk perception
8. Objective risk exists
9. Risk is determined by the historical data
10. Risk relates to negative consequences
11. Risk and probability cannot be determined in case of large uncertainties
12. There are large inherent uncertainties in risk analyses
13. Risk acceptance criteria contribute to obtaining a high safety standard
14. ALARP can be verified by cost-benefit analyses
15. The cautionary/pre-cautionary principles and risk management cannot be meaningfully integrated.

2.1 Risk is equal to the expected value

It is common to refer to risk as probability multiplied by consequences (losses), i.e. what is called the expected value in probability calculus. If C is the quantity of interest, for example the number of future attacks, the number of fatalities, the costs etc., the expected value would be a good representation of risk if this value is approximately equal to C , i.e. $EC \approx C$. But since C is unknown at the time of the assessment, how can we be sure that this approximation would be accurate? Can the law of large numbers be applied, expressing that the empirical mean of independent identically distributed random variables converges to the expected value when the number of variables increases to infinity? Or the portfolio theory (Levy & Sarnat 1990) saying that the value of a portfolio of projects is approximately equal to the expected value, plus the systematic risk (uncertainties) caused by events affecting the whole market?

Yes, it is likely that if C is the sum of a number of projects, or some average number, our expected value could be a good prediction of C . Take for example the number of fatalities in traffic in a specific country. From previous years we have data that can be used to accurately predict the number of fatalities next year (C). In Norway about 250 people were killed last year, and using this number as EC and predictor for the coming year, we would be quite sure that this number is close to the actual C .

However, in many cases the uncertainties are much larger. Looking at the number of fatalities in Norway caused by terrorist attacks the next year, the historical data would give a poor basis. We may assign an EC but obviously EC could be far away from C . The accuracy increases when we extend the population of interest. If we look at one unit (e.g. country) in isolation the C number is in general more uncertain than if we consider many units (e.g. countries). Yet, there will always be uncertainties, and in a world where the speed of change is increasing, relevant historical data are scarce and will not be sufficient to obtain accurate predictions.

Even so, some researchers define risk by the expected value. Consider the terrorism case discussed in Willis (2007). Willis (2007) defines risk as follows:

Terrorism risk: The expected consequences of an existent threat, which for a given target, attack mode, target vulnerability, and damage type, can be expressed as

$$\text{Risk} = P(\text{attack occurs}) \cdot P(\text{attacks results in damage} \mid \text{attacks occurs}) \cdot E[\text{damage} \mid \text{attacks occurs and results in damage}]$$

Willis (2007) refers to Haimes (2004) who highlights that expected value decision-making is misleading for rare and extreme events. The expected value (the mean or the central tendency) does not adequately capture events with low probabilities and high consequences. Nonetheless, Willis represents risk by the expected value as the basis for his analysis. The motivation seems to be that the expected value provides a suitable practical approach for comparing and aggregating terrorism risk, as it is based on just one number.

For terrorism risk, where the possible consequences could be extreme and the uncertainties in underlying phenomena and processes are so large, it is however obvious that the expected value may hide important aspects of concern for risk management. The expected value can be small, say 0.01 fatalities, but extreme events with millions of fatalities may occur, and this needs special attention.

Hence we need to see beyond the expected values. We have to take into account uncertainties and risks. Risk management is concerned about how to assess these uncertainties and risk, and how to handle them.

2.2 Risk equals uncertainty

Risk is sometimes associated with uncertainty, for example, in Cabinet Office (2002), risk refers to uncertainty of outcome, of actions and events. Often the uncertainty is seen in relation to the expected value, and the variance is used as a measure of risk.

As an example, consider the problem of investing money in a stock market. Suppose the investor considers two alternatives, both with expectation 1, and variances 0.16 and 0.08, respectively. As alternative 2 has the lowest risk (uncertainty), expressed by the variance, this alternative would normally be chosen.

As another example, consider the number of fatalities in traffic next year in a specific country. Then the variance is rather small, as the number of fatalities shows rather small variations from year to year. Hence according to this definition of risk, we must conclude that the risk is small, even

though the number of fatalities are many thousands each year. Clearly, this definition of risk fails to capture an essential aspect, the consequence dimension. Uncertainty cannot be isolated from the intensity, size, extension etc. of the consequences. Take an extreme case where only two outcomes are possible, 0 and 1, corresponding to 0 and 1 fatality, and the decision alternatives are A and B, having uncertainty (probability) distributions (0.5,0.5), and (0.0001, 0.9999), respectively. Hence for alternative A there is a higher degree of uncertainty than for alternative B, meaning that risk according to this definition is higher for alternative A than for B. However, considering both dimensions, both uncertainty and the consequences, we would of course judge alternative B to have the highest risk as the negative outcome 1 is nearly certain to occur.

2.3 Risk is equal to an event

Risk is also being defined as an event or as a consequence:

1. Risk is a situation or event where something of human value (including humans themselves) is at stake and where the outcome is uncertain (Rosa 1998, 2003)
2. Risk is an uncertain consequence of an event or an activity with respect to something that human value (Renn 2005).

We do not distinguish between the definitions 1) and 2). They are basically expressing the same: Risk is an event or a consequence of an event. See Figure 1. The activity considered could produce events and consequences and these are subject to uncertainties. Something of human value is at stake.




Figure 1. Risk defined as an event or a consequence (Aven & Renn 2008)

According to these definitions, risk expresses a state of the world independent of our knowledge and perceptions. Referring to risk as an event or a consequence, we cannot conclude on risk being high or low, or compare options with respect to risk. Compared to standard terminology in risk research and risk management, they lead to conceptual difficulties that are incompatible with the everyday use of risk in most applications, as discussed in the following.

An analogous definition to 1-2) is found in reliability theory. Here the term unavailability is normally used as the expected fraction of time the system being considered is unavailable, i.e. is not functioning (Aven & Jensen 1999), but we also see unavailability defined in the sense 1 as a state of the world, expressed by the actual fraction of time the system is unavailable (ISO 2005). Then we may consider failures in the system as the “events” according to the definition 1 and the fractional downtime as the consequences. The events and consequences are subject to uncertainties.

The definitions 1-2) mean that risk and assessment of risk are separate domains of the world. The occurrence of a leakage in a process plant is a risk (according to 1). This event is subject to

uncertainties, but the risk concept is restricted to the event “leakage” – the uncertainties and how people judge the uncertainties is a different domain. Hence a risk assessment according to 1-2) cannot conclude for example that the risk is high or low, or that option A has a lower or higher risk than option B, as it makes no sense to speak about a high or higher leakage. Instead the assessment needs to conclude on the uncertainty or the probability of the risk being high or higher.

A similar comment can be made on risk perception.

We conclude that any judgement about risk need to take into account uncertainties/likelihoods, so why not include this dimension into the risk concept?

We refer to Aven and Renn (2008) for further discussion of this risk perspective.

2.4 Risk is a probability or a probability distribution

We often see risk defined by probabilities. Here are some examples:

- Risk is the probability of an undesirable event (Campbell 2005)
- Risk is the probability of an adverse outcome (Graham & Weiner 1995)
- Risk is a measure of the probability and severity of adverse effects (Lowrance 1976)
- Risk is the combination of probability of an event and its consequences (ISO 2002)
- Risk is defined as a set of scenarios s_i , each of which has a probability p_i and a consequence c_i (Kaplan & Garrick 1981, Kaplan 1991).

Clearly, the first definition is inadequate as a description of risk, as the consequences and outcomes are not taken into account. If we consider the undesirable event “machine failure”, the consequences can range from negligible to disaster depending on the availability and performance of a set of barriers, as well as the extent of exposure of human lives, and other objects that human values.

However, also the other probability based definitions can be challenged. A probability is not capturing all aspects of concern. To explain this we need to first introduce the two common ways of interpreting a probability: the classical relative frequency interpretation and the subjective Bayesian interpretation.

According to the classical relative frequency paradigm, a probability is interpreted as the relative fraction of times the events occur if the situation analyzed were hypothetically “repeated” an infinite number of times. The underlying probability is unknown, and is estimated in the risk analysis. Hence if this interpretation is adopted in the above definitions of risk, we have to take into account that the risk estimates could be more or less accurate relative to the underlying true risk. The uncertainties in the estimates could be very large, and difficult to express.

The alternative (the Bayesian perspective) considers probability as a measure of uncertainty about events and outcomes (consequences), seen through the eyes of the assessor and based on the available background information and knowledge. Probability is a subjective measure of uncertainty, conditional on the background information. The reference is a certain standard such as drawing a ball from an urn. If we assign a probability of 0.4 for an event A , we compare our uncertainty of A to occur with drawing a red ball from an urn having 10 balls where 4 are red. Objective probabilities do not exist.

Following this perspective, we assign a probability by performing uncertainty assessments, and there is no reference to a correct or true probability. A probability is always conditional on the background knowledge, and given this background knowledge there is no uncertainty related to the assigned probability, as it is an expression of uncertainty.

However, a probability is not a perfect tool for this purpose. The assigned probabilities are conditional on a specific background knowledge, and they could produce poor predictions.

Surprises relative to the assigned probabilities may occur, and by just addressing probabilities such surprises may be overlooked (Aven 2007a, 2008).

Let us look at two examples:

Maintenance

We consider an offshore petroleum installation where the operations management is concerned about the deterioration of some critical equipment. The maintenance discipline ensures that the deterioration will not cause safety problems. It refers to a special maintenance program that will be implemented, which will cope with the problem. So what is the risk associated with hydrocarbon leakages caused by operational problems. Given the background information of the maintenance discipline, a leakage probability (for a defined size) of 10% is assigned. This number is based on relevant historical data, and do not in any respect reflect the concern of the operations management. The assignment assumes that the maintenance program will be effective. But surprises could occur. Production of oil over time lead to changes in operating conditions, such as increased production of water, H₂S and CO₂ content, scaling, bacteria growth, emulsions, etc.; problems that to large extent need to be solved by the addition of chemicals. These are all factors causing increased likelihood of corrosion, material brittleness and other conditions that may cause leakages.

By the assignment of 10% we hide an important element of uncertainty. In a risk analysis a number of such probability assignments are performed, and the hidden uncertainties could create surprising outcomes someplace. You do not know where it will come, but it certainly could happen.

Offshore diving activities

Consider the risk, seen through the eyes of a risk analyst in the 1970s, related to future health problems for divers working on offshore petroleum projects. An assignment is to be made for the probability that a diver would experience health problems (properly defined) during the coming 30 years due to the diving activities. Let us assume that an assignment of 1% is made. This number is based on the available knowledge at that time. There are not strong indications that the divers will experience health problems. However, we know today, that these probabilities led to poor predictions. Many divers have experienced severe health problems (Aven & Vinnem 2007, p. 7). By restricting risk to the probability assignments alone, we see that aspects of uncertainty and risk are hidden. There is a lack of understanding about the underlying phenomena, but the probability assignments alone are not able to fully describe this status.

2.5 Risk equals expected disutility

If X is the outcomes (consequences) and $u(X)$ the utility function, risk defined by the expected disutility is given by $-Eu(X)$ (Campbell 2005). Hence the preferences of the decision maker is a part of the risk concept. The result is a mixture of scientific assessments of uncertainties about X and the decision makers preferences concerning different values of X . We consider this to be an unfortunate mixture. There will be a strong degree of arbitrariness in the choice of the utility function, and some decision makers would also be reluctant to specify the utility function as it reduces their flexibility to weight different concerns in specific cases. Risk should be possible to describe also in case that the decision maker is not able or willing to define his/her utility function.

2.6 Risk is restricted to the case of known probabilities

In economic applications a distinction has traditionally been made between risk and uncertainty, based on the availability of information. Under risk the probability distribution of the

performance measures can be assigned objectively, whereas under uncertainty these probabilities must be assigned or estimated on a subjective basis (Douglas 1983). This definition goes back to Knight (1921).

Although this definition is often referred to, it is not so often used in practice. The problem is of course that we seldom have known distributions, and then we cannot refer to the risk concept. The Knightian definition violates the intuitive interpretation of risk, which is related to situations of uncertainty and lack of predictability.

2.7 Risk based on subjective probabilities is the same as risk perception

In cultural theory and constructivism, *risk is the same as risk perception* (Rosa 1998). Risk coincides with the perceptions of it (Douglas & Wildavsky 1982, Freudenburg 1989). Beck (1992) concludes that “because risks are risks in knowledge, perceptions of risks and risk are not different things, but one and the same”. Beck argues that the distinction between risk and risk perception is central to a scientific myth of expertise, according to which the population “perceives risks” but science determines (i.e., identifies and quantifies) risk (Campbell & Currie 2006, p. 152).

This viewpoint of risk being the same as risk perception is, however, not confined to these paradigms and scientists (Rosa 1998). Rosa (1998) refers for example to the leading risk psychometrician Paul Slovic who has written: ‘Human beings have invented the concept of “risk” . . . there is no such thing as “real risk” or “objective risk” (Slovic 1992: 119).

But rejecting the idea that there exists a “real risk” or an “objective risk”, does not mean that risk is the same as perceived risk. If probability is a way of expressing uncertainties, seen through the eyes of the assigner (a subjective probability), there is no “real risk” or “objective risk”. However, subjective probabilities and related risk assignments are not the same as risk perception. You may assign a probability equal to 0.000000001 for an event to occur, but still find the risk to be intolerable. Our judgments about risk are as we know from many risk perception studies influenced by a number of factors outside the realm of the probabilistic world. The assigned probability and the risk perceptions are different dimensions, or separate domains of the world using Rosa’s words.

In the case that the risk perspective is based on the idea that a true risk exists, it is obvious that the risk = risk perception thesis is wrong, refer Campbell (2005, p. 230). The above analysis shows that this thesis is invalid also for other risk perspectives.

We refer to Aven and Renn (2008).

2.8 Objective risk exists

The classical relative frequency approach to risk is based on the idea of an underlying true, objective probability. We understand the meaning of this perspective in gambling situations, but what does this idea of a true probability mean in a more complex context. Consider for example the probability of a terrorist attack, i.e. $P(\text{attack occurs})$. How can this probability be understood by reference to a thought-constructed repeated experiment?

It does not work at all. It makes no sense to define a large set of “identical”, independent attack situations, where some aspects (for example related to the potential attackers and the political context) are fixed and others (for example the attackers motivation) are subject to variation. Say that the attack probability is 10%. Then in 1000 situations, with the attackers and the political context specified, the attackers will 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 random situations such “experiments” are meaningful, but not for deliberate actions as in the terrorism case.

As another example, consider the probability of at least one fatality during one year in a production facility. According to the relative frequency view, this probability is interpreted as the proportion of facilities with at least one fatality when considering an infinite number of similar facilities. This is of course a thought experiment - in real life we just have one such facility. Therefore, the probability is not a property of the unit itself, but the population it belongs to. How should we then understand the meaning of similar facilities? Does it mean the same type of buildings and equipment, the same operational procedures, the same type of personnel positions, the same type of training programmes, the same organizational philosophy, the same influence of exogenous factors, etc. As long as we speak about similarities on a macro level, the answer is yes. But something must be different, because otherwise we would get exactly the same output result for each facility, either the occurrence of at least one fatality or no such occurrence. There must be some variation on a micro level to produce the variation of the output result. So we should allow for variations in the equipment quality, human behaviour, etc. But the question is to what extent we should allow for such variation. For example, in human behaviour, do we specify the safety culture or the standard of the private lives of the personnel, or are these factors to be regarded as factors creating the variations from one facility to another, i.e. the stochastic (aleatory) uncertainty? We see that we will have a hard time specifying what should be the framework conditions of the experiment and what should be stochastic uncertainty. In practice we seldom see such a specification carried out, because the framework conditions of the experiment are tacitly understood. As seen from the above example, it is not obvious how to make a proper definition of the population.

2.9 Risk is determined by the historical data

To many people, risk is closely related to accident statistics. Numerous reports and tables are produced showing the number of fatalities and injuries as a result of accidents. The statistics may cover the total number of accidents associated with an activity within different consequence categories (loss of lives, personal injuries, material losses, etc.) and they could be related to different types of accidents, such as industrial accidents and transport accidents.

Often the statistics are related to time periods, and then time trends can be identified. More detailed information is also available in some cases, related to, for example, occupation, sex, age, operations, type of injury, etc.

Do these data provide information about the future, about risk? Yes, although the data are historical data, they would usually provide a good picture of what to expect in the future. If the numbers of accidental deaths in traffic during the previous five years are 1000, 800, 700, 800, 750, we know a lot about risk, even though we have not explicitly expressed it by formulating predictions and uncertainties. This is risk related to the total activity, not to individuals. Depending on your driving habits, these records could be more or less representative for you.

However, historical data may exclude extreme observations, but this does not preclude such observations to occur in the future. The historical figures can obviously lead to poor predictions.

By attempting to understand the data, by looking for trends, we may be able to improve the predictions. But we may also end up “over-interpreting” the data in the sense that we look for all sorts of explanations for why the historical figures are as they are. The data may indicate that the quantity of interest (for example the injury rate) is increasing, but perhaps the trend arrow will be reversed next month. We can analyse possible underlying conditions that can affect the quantity, but it is not easy to reflect what the important factors are, and what is “noise” or arbitrariness.

An analysis based on the historical numbers could easily become too narrow and imply that extreme outcomes are ignored. Surprises occur from time to time, and suddenly an event could occur that dramatically changes the development, with the consequence that the quantity of interest

jump up or down. In a risk analysis such events should be identified. However, the problem is that we do not always have the knowledge and insights to be able to identify such events, because they are extremely unexpected.

2.10 Risk relates to negative consequences

Most people associate the word risk with something undesirable and negative. There are, however, good reasons for not restricting the risk concept to negative consequences, and many definitions of risk relate risk to both negative and positive consequences. What is a negative consequence or outcome? To some, an outcome can be negative, and for others a positive outcome. We wish to avoid a discussion on whether a consequence is classified in the correct category. In a risk assessment, the aim is to uncover all relevant consequences, then assess uncertainties and assign probabilities.

2.11 Risk and probability cannot be determined in case of large uncertainties

It is common to hear statements saying that risk and probability cannot be determined in case of large uncertainties. It is however a myth. Risk analyses can always be carried out. Risk can always be expressed, regardless of access to input data. Through the risk analysis the knowledge and lack of knowledge one has concerning various quantities are expressed. Of course, in a case of large uncertainties, it will be difficult to establish good predictions, but the purpose of the analysis is to describe the uncertainties. See the following section.

2.12 There are large inherent uncertainties in risk analyses

The risk analysis systemizes available knowledge and uncertainties about phenomena, systems, and activities that are being studied. What can go wrong, why, and what are the consequences? This knowledge and this uncertainty are described and discussed, and thereby we obtain a basis on which we can evaluate what is important and compare different solutions and measures.

If one has a large and relevant data base, the probabilities derived from it could be precise in the sense that they may be able to provide accurate predictions of future events. For example, assume that one has observed 200 failures in a population of 10 000 units of type T over a one year period. The derived probability of failure for one arbitrary chosen unit is then 2%, and we will predict for example 20 failures per thousand units. We can express the uncertainty, for example, using a 95% prediction interval: [13, 31]. The number of failures will lie within this interval with a 95% probability. To establish this interval, let X denote the number of failures among 1000 units. Then X has a binomial distribution, which can be approximated by a normal distribution with mean 20 and standard deviation 4.4, and this gives $P(11 \leq X \leq 29) = 0.95$.

In a risk analysis context, we often focus on rare events, for example, the occurrence of a fatal accident, an accident that causes impairment of a main safety function, etc. We have only one unit or activity, and we are able to give a good prediction about the future: no fatal accidents will occur the next year. Fortunately, such a prediction will normally provide correct results. The risk analysis however, should also express the likelihood associated with whether the event will occur. This raises the question about precision in the probability assignment.

Many risk analyses today are characterized either by silence on the subject, or by general statements such as:

The analyses are based on the “best estimates” obtained by using the company’s standards for models and data. It is acknowledged that there are uncertainties associated with all elements in the analysis, from the hazard identification to the models and probability calculations. It is concluded that the precision of the analysis is limited, and that one must take this into considerations when comparing the results with the risk acceptance criteria and tolerability limits.

The above statement is not very convincing, and it is obvious that there is no clarity regarding what the analyses express, and what uncertainty means in a risk analysis context.

In any event, does this acknowledgment -- that a considerable amount of uncertainty exists -- affect the analyses and the conclusions? Only very rarely! Our impression is that one writes such statements just to meet a requirement, and then they are put aside. This says a lot about the quality of the analyses.

If the goal of the risk analysis is to obtain reliable, i.e. accurate, estimates of some true risk, we can quickly conclude that risk analysis fails as a scientific method. Referring to Section 2.4, we can conclude that the classical approach to risk analysis does not work in situations of large uncertainties. The uncertainties of the risk estimates are too large.

Alternatively, we may consider risk analysis as a tool for assessing uncertainties about risk and risk estimates. Since the assessment's aim then is to express uncertainties about the true risk, reliability is not related to the accuracy in the results but rather the “accuracy” of the transformation of uncertainties to probabilities. Risk analysis is then not about bounding and reducing uncertainties, but to describe uncertainties. Two prevailing approaches for describing the uncertainties are:

Traditional statistical methods such as confidence intervals

The probability of frequency approach, i.e. assessing epistemic uncertainties about the risk by means of subjective probabilities. In this approach there are two levels of probability introduced; i) the relative frequency interpreted probabilities reflecting variation within populations and ii) the subjective probabilities reflecting the analyst's uncertainty what the correct relative frequency probabilities are (see e.g. Kaplan & Garrick (1981) and Aven (2003)). In Garrick et al. (2004) the probability of frequency approach is suggested for risk analysis of attacks. Garrick et al. (2004) refer to a probability distribution saying for example that there is a probability of 20% that the attackers would succeed in 10% of their attacks.

However, confidence intervals would not work in this setting as we do not have sufficient amount of relevant data.

The ambition of the probability of frequency approach is to express the epistemic uncertainties of the probability p of an attack, and take into account all relevant factors causing uncertainties. The analysis may produce a 90% credibility interval for p , $[a, b]$, saying that the analyst is 90% confident that p lies in the interval $[a, b]$. In practice it is difficult to perform a complete uncertainty analysis following this approach. In theory an uncertainty distribution on the total model and parameter space should be established, which is impossible to do. So in applications only a few marginal distributions on some selected parameters are normally specified, and therefore the uncertainty distributions on the output probabilities are just reflecting some aspects of the uncertainty. This makes it difficult to interpret the produced uncertainties.

If the risk perspective is based on probability being a measure of uncertainty seen through the eyes of the assessor, and based on a background knowledge, we can argue along the same lines. As for the probability of frequency approach, we conclude that this approach in general meets the reliability requirement, if reliability is associated with subjective probability assignments and these follow the standards established for such assignments.

2.13 Risk acceptance criteria contribute to obtaining a high standard of safety

To manage safety and security, it is common to use a hierarchy of goals, criteria and requirements, such as

- overall ideal goals, for example “our goal is to have no accidents”
- risk acceptance criteria (defined as upper limits of acceptable risk) or tolerability limits, controlling the accident risk, for example “the individual probability of being killed in an accident shall not exceed 0.1 %”
- requirements related to the performance of safety systems and barriers, such as a reliability requirement for a safety system
- requirements related to the specific design and operation of a component or subsystem, for example the gas detection system.

According to the standard procedures for using such goals, criteria and requirements, they are to be specified before alternatives are generated and subsequently analysed. The point is to look for what to obtain before looking for possible ways of implementation. For example, the Norwegian offshore petroleum regulations state that risk acceptance criteria (expressed as upper limits of acceptable risk) should be developed, and before the risk analyses are carried out (PSA 2001, Aven 2007b). Note that we in the following, when using the term “risk acceptance criteria”, always have in mind such upper limits.

Are such criteria appropriate for managing investments in safety and security? With large uncertainties, it is not meaningful to use such limits as the precision level is poor. However, it is also questionable to use such criteria where there is no discussion about the risk picture as such.

Consider the following criterion for an offshore installation:

The probability of getting an oil spill during one year of operation causing an environmental damage having a restitution period of more than z years, should not exceed $1 \cdot 10^{-x}$.

At the political level it is obvious that it would not be possible to establish consensus about such a limit. Different parties, would have different preferences. But for the Government it should be possible to establish such a number? Say that it would make an attempt to do this. And suppose that it considers two options, a weak limit, say $1 \cdot 10^{-3}$ and a strong limit say $1 \cdot 10^{-4}$. What limit should it choose? The answer would be the weak limit, as the strong limit could mean lack of flexibility in choosing the overall best solution. If the benefits are sufficient large, the level $1 \cdot 10^{-3}$ could be acceptable. Following this line of arguments, the use of such limits leads to the formulation of weak limits, which are met in most situations. Risk assessments are then used to test whether risks are acceptable in relation to these weak limits. It is to large extent waste of money, the conclusions are obvious.

At the operational level, the same type of arguments will apply. The oil company is to determine an acceptance criterion, and it faces the same type of dilemmas as above. Why should it specify strong limits? It would restrict the company from obtaining the overall best solutions. The result is that weak limits are specified and risk assessments play the role of verification, a role that adds not much value.

If a high level of safety or security is to be obtained, other mechanisms need to be implemented than risk acceptance criteria. If such criteria are established, they give a focus on obtaining a minimum safety standard, instead of continuous improvement and risk reduction.

The ALARP principle represents such a mechanism. The ALARP principle expresses that the risk should be reduced to a level that is as low as reasonably practicable. A risk reducing measure

should be implemented provided it cannot be demonstrated that the costs are grossly disproportionate relative to the gains obtained (HSE 2001).

2.14 ALARP can be verified by cost-benefit analyses

The ALARP principle implies what could be referred to as the principle of ‘reversed onus of proof’. This implies that the base case is that all identified risk reduction measures should be implemented, unless it can be demonstrated that there is gross disproportion between costs and benefits. To verify ALARP, procedures mainly based on engineering judgments and codes are used, but also traditional cost-benefit analyses and cost effectiveness analyses. When using such analyses, guidance values are often used, to specify what values that define ‘gross disproportion’.

The practice of using traditional cost-benefit analyses and cost effectiveness analyses to verify ALARP has been questioned (Aven & Abrahamsen 2007). The ALARP principle is an example of application of the cautionary principle (see Section 2.15). Uncertainty should be given strong weight, and the grossly disproportionate criterion is a way of making the principle operational. However, cost-benefit analyses calculating expected net present values to large extent ignore the risks (uncertainties) and the use of this approach to weight the risk and uncertainties is therefore problematic. The same applies to the cost-effectiveness indices such as the expected cost per expected number of saved lives (referred to as the implied cost of averting a statistical fatality, ICAF) which are often used instead of full cost-benefit analyses. If a measure costs 2 million euros and the risk analysis shows that the measure will bring about a reduction in the number of expected fatalities by 0.1, then the ICAF is equal to $2/0.1 = 20$ million euros. By comparing this number with reference values, we can express the effectiveness of the measure.

Modifications of the traditional cost-benefit analysis are suggested to cope for this problem, see e.g. Aven & Flage (2008). In these methods, adjustments are made on either the discount rate or the contribution from the cash flows. This latter case could be based on the use of certainty equivalents for the uncertain cash flows. Although arguments are provided to support these methods, their rationale can be questioned. There is a significant element of arbitrariness associated with the methods, in particular when seen in relation to the standard given by the expected utility theory.

2.15 The cautionary/pre-cautionary principles and risk management cannot be meaningfully integrated

It is common among engineers, economist and others to consider a rational decision in face of uncertainties to be based on risk and decision analysis, including cost-benefit analysis. It is irrational to give weight to the cautionary and pre-cautionary principles. However, a further look into this view shows that is cannot be justified.

The cautionary principle is a basic principle in risk management, expressing that in the face of uncertainty, *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 & Vinnem 2007). The level of caution adopted will of course have to be balanced against other concerns such as costs. However, all industries would introduce some minimum requirements to protect people and the environment, and these requirements can be considered justified by the reference to the cautionary principle.

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

which 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 that fires occur from time to time in such plants. It does not matter whether we calculate a fire probability of x or y , as long as we consider the risks to be significant; and this type of risk has been judged to be significant by the authorities. The justification is experience from similar plants and sound judgments. A fire may occur, since it is not an unlikely event, and we should then be prepared. We need no references to cost-benefit analysis. The requirement is based on a cautionary thinking.

Risk analyses, cost-benefit analyses and similar types of analyses are tools providing insights into risks and the trade-offs involved. But they are just tools - with strong limitations. Their results are conditioned on a number of assumptions and suppositions. The analyses do not express objective results. Being cautious also means reflecting this fact. We should not put more emphasis on the predictions and assessments of the analyses than what can be justified by the methods being used.

In the face of uncertainties related to the possible occurrences of failures, hazardous situations and accidents, we are cautious and adopt principles of risk management, such as

- robust design solutions, such that deviations from normal conditions are not leading to hazardous situations and accidents,
- design for flexibility, meaning that it is possible to utilise a new situation and adapt to changes in the frame conditions,
- implementation of safety barriers, to reduce the negative consequences of hazardous situations if they should occur, for example a fire,
- improvement of the performance of barriers by using redundancy, maintenance/testing, etc.
- quality control/ quality assurance,
- the precautionary principle, saying that in the case of lack of scientific certainty on the possible consequences of an activity, we should not carry out the activity.
- the ALARP-principle, saying that the risk should be reduced to a level which is as low as reasonably practicable.

Thus the precautionary principle may be considered a special case of the cautionary principle, as it is applicable in cases of scientific uncertainties (Sandin 1999, Löfstedt 2003, Aven 2006). There are however many definitions of the precautionary principle. The well-known 1992 Rio Declaration use the following definition:

In order to protect the environment, the precautionary approach shall be widely applied by States according to their capabilities. Where there are threats of serious or irreversible damage, lack of full scientific certainty shall not be used as a reason for postponing cost-effective measures to prevent environmental degradation.

Seeing beyond environmental protection, a definition such as the following reflects what is a typical way of understanding this principle:

The precautionary principle is the ethical principle that if the consequences of an action, especially the use of technology, are subject to scientific uncertainty, then it is better not to carry out the action rather than risk the uncertain, but possibly very negative, consequences.

In the following we will refer to the cautionary and precautionary principles, and in this way avoid a discussion about whether we refer to the cautionary principle or the precautionary principle. The distinction is not essential for the purpose of this paper.

We have to acknowledge that there exists no simple and mechanistic method or procedure for dealing with uncertainties and balancing different concerns. This is also recognised by many others

analysts, see e.g. the approach adopted by the risk governance framework (Renn 2005) and the risk framework used by the UK Cabinet Office (2002).

Uncertainty is an important aspect of risk, and hence the cautionary and precautionary principles constitute important aspect of risk management.

3. CONCLUSIONS

We suggest the following definition of risk (Aven 2007a):

By risk we understand the two-dimensional combination of i) events A and the consequences of these events C, and ii) the associated uncertainties U (wil A occurs and what value will C take) (I)

We refer to this definition as the (C,U) risk definition. For simplicity, we write only C, instead of A and C.

We may rephrase this definition by saying that risk associated with an activity is to be understood as (Aven & Renn 2008):

Uncertainty about and severity of the consequences of an activity (I'),

where severity refers to intensity, size, extension, and so on, and is with respect to something that humans value (lives, the environment, money, etc). Losses and gains, for example expressed by money or the number of fatalities, are ways of defining the severity of the consequences.

The main features of the definition are illustrated in Figure 2.




Figure 2. Illustration of the risk definition (I, I')

The uncertainty relates to both the event and the consequences given that this event occurs.

We see that the definition is based on the combined dimensions Consequences (outcome stakes) and Uncertainties. It is essential that the second dimension is uncertainties and not probabilities. The concept of risk should cover uncertainties beyond probabilities. Probabilities and expected values are not perfect tools for expressing uncertainties, as discussed in Section 2.4.

Note that following our definition of risk (I), a low degree of uncertainty does not necessarily mean a low risk, or a high degree of uncertainty does not necessarily mean a high level of risk. As risk is defined as the two-dimensional combination of consequences and uncertainties, any judgment about the level of risk, needs to consider both dimensions. See example in Section 2.2.

The risk concept is supporting a broad perspective on risk, as for example shown by the way risk assessment are conducted. In stead of a traditional quantitative risk assessment approach, we

recommend a more qualitative approach or semi-quantitative approach (Aven 2008). The basic features of the approach can be summarised as follows:

A broad qualitative risk picture is established highlighting

- Potential hazards/threats and accident scenarios
- Barriers and the effectiveness of these barriers
- Risk influencing factors and possible risk reducing measures
- Uncertainties in phenomena and processes
- Vulnerabilities
- Special features of the consequences
- Manageability factors

Crude risk categorisations are defined based on this risk picture, reflecting

- Probabilities/frequencies of hazards/threats
- Expected losses given the occurrence of such a hazard/threat
- Factors that could create large deviations between expected outcomes and the actual outcomes (uncertainties, vulnerabilities)

Evaluations of the risk picture and categorisations to compare alternatives and make judgments about risk acceptance.

Quantifying risk using risk indices such as expected number of fatalities gives an impression that risk can be expressed in a very precise way. However, in most cases, the arbitrariness is large, and the semi-quantitative approach acknowledges this by providing crude risk numbers, including assessments of the factors that can cause “surprises” relative to the probabilities and expected values. We are not opposed to detailed risk quantification as such, but quantification often requires strong simplifications and assumptions and as result, important factors could be ignored or given too little (or much) weight. In a qualitative or semi-quantitative analysis a more comprehensive risk picture can be established, taking into account underlying factors influencing risks. In contrast to the prevailing use of quantitative risk assessments, the precision level of the risk description is in line with the accuracy of the risk assessment tool. In addition, risk quantification is very resource demanding. We need to ask whether the resources are used in the best way. We conclude that in many cases more is gained by pursuing a broader more qualitative approach, which allows for considerations beyond the probabilities and expected values.

For problems with large uncertainties, risk assessments could support decision making, but other principles, measures and instruments are required, such as the cautionary principle.

References

- Apostolakis, G.E. and Lemon, D.M. 2005. A Screening Methodology for the Identification and Ranking of Infrastructure Vulnerabilities Due to Terrorism. *Risk Analysis*, 24: 361-376.
- AS/NZS 4360 (2004) Australian/New Zealand Standard: Risk management.
- Aven, T. 2003. Foundations of Risk Analysis – A Knowledge and Decision Oriented Perspective, N.Y.: Wiley.
- Aven, T. 2006. On the precautionary principle, in the context of different perspectives on risk. *Risk Management: an International Journal*, 8: 192-205.
- Aven T. 2007a. A unified framework for risk and vulnerability analysis and management covering both safety and security. *Reliability Engineering and System Safety*, 92: 745-754.
- Aven, T. 2007b. On the ethical justification for the use of risk acceptance criteria. *Risk Analysis*, 27: 303-312.
- Aven, T. 2008. A semi-quantitative approach to risk analysis, as an alternative to QRAs. *Reliability Engineering & Systems Safety*, 93: 768-775.

- Aven, T. and Abrahamsen, E.B. 2007. On the use of cost-benefit analysis in ALARP processes. *I. J. of Performability*, 3: 345-353.
- Aven, T. and Flage, R. 2008. Use of decision criteria based on expected values to support decision making in a production assurance and safety setting. *Reliability Engineering & Systems Safety*. To appear.
- Aven, T. and Jensen, U. 1999. *Stochastic Models in Reliability*. N.Y.: Springer.
- Aven, T. and Kristensen, V. 2005. Perspectives on risk – Review and discussion of the basis for establishing a unified and holistic approach. *Reliability Engineering and System Safety*, 90, 1-14.
- Aven, T. and Renn, O. 2008. On risk defined as an event where the outcome is uncertain. *J. Risk Research*. To appear.
- Aven, T., Vinnem, J.E. and W. Røed 2006. On the use of goals, quantitative criteria and requirements in safety management. *Risk Management: an International Journal*, 8: 118-132.
- Aven, T. and Vinnem, J.E. 2005. On the use of risk acceptance criteria in the offshore oil and gas industry. *Reliability Engineering and System Safety*, 90: 15-24.
- Aven, T. and Vinnem, J.E. 2007. *Risk Management, with Applications from the Offshore Oil and Gas Industry*. N.Y.: Springer Verlag.
- Barlow, R.E. 1998. *Engineering Reliability*. Pennsylvania: Siam – Society for Industrial and Applied Mathematics.
- Beck, U. 1992. *Risk Society: Toward a New Modernity*, translated by Mark A. Ritter, London: Sage Publications.
- Bedford, T. and Cooke, R. 2001. *Probabilistic Risk Analysis*. Cambridge: Cambridge U. Press.
- Cabinet Office 2002. Risk: improving government's capability to handle risk and uncertainty. Strategy unit report. UK.
- Campbell, S. 2005. Determining overall risk. *J. of Risk Research*, 8: 569-581.
- Campbell, S. and Currie, G. 2006. Against Beck. In defence of risk analysis. *Philosophy of the Social Sciences*, 36: 149-172.
- Douglas, E.J. 1983. *Managerial Economics: Theory, Practice and Problems*. 2nd ed., N.J.: Prentice Hall.
- Douglas, M. and Wildavsky, A. 1982. Risk and Culture: The Selection of Technological and Environmental Dangers, Berkeley, CA: University of California Press.
- Fischhoff, B. 1994. Acceptable risk: A conceptual proposal. *Risk: Health, Safety & Environment*, 1: 1-28.
- Freudenburg, W.R. 1989. Perceived risk, real risk: social science and the art of probabilistic risk assessment, *Science* 242, 44-49.
- Garrick, J.B., et al. 2004. Confronting the risks of terrorism: making the right decisions. *Reliability Engineering and System Safety*, 86: 129-176.
- Graham, J.D. and Weiner, J.B. (eds) 1995. *Risk versus Risk: tradeoffs*. Protecting Health and the Environment, Cambridge: Harvard University Press.
- Guikema, S.D. 2008. Modeling Intelligent Actors in Reliability Analysis: An Overview of the State of the Art, in Bier, V.M. and N. Azaiez, Eds. *Combining Reliability and Game Theory*, Springer Series on Reliability Engineering. In press.
- Guikema, S.D. and Aven, T. 2008. Assessing Risk from Intelligent Attacks: A Perspective on Approaches. Submitted for publication.
- Haimes, Y.Y. 2004. *Risk Modelling, Assessment, and Management*, 2nd ed. N.J.: Wiley.
- HES 2001. Reducing risk, protecting people. HES Books, ISBN 0 71762151 0.
- HSE 2006. Offshore installations (safety case) regulations 2005 regulation 12 demonstrating compliance with the relevant statutory provisions
- ISO 2002. Risk management vocabulary. ISO/IEC Guide 73.
- ISO 2005. Petroleum, petrochemical and natural gas industries—Production assurance and reliability management. CD 20815. Committee draft, 2005.
- Kaplan, S. and B.J. Garrick, 1981. On the quantitative definition of risk. *Risk Analysis*, 1: 11-27.
- Kaplan, S. 1991. Risk Assessment and Risk Management – Basic Concepts and Terminology. In *Risk Management: Expanding Horizons in Nuclear Power and Other Industries*, Hemisphere Publ. Corp., Boston, MA, 11-28.
- Kasperson, R.; Renn, O.; Slovic P.; Brown, H.; Emel, J.; Goble, R.; Kasperson, J.X. and Ratick, S. 1988. The Social Amplification of Risk. A Conceptual Framework. *Risk Analysis*, 8 : 177-187.
- Kasperson, J.X.; Kasperson, R.E.; Pidgeon, N. and Slovic, P. 2003. The Social Amplification of Risk: Assessing Fifteen Years of Research and Theory, in: N. Pidgeon; R.E. Kasperson and P. Slovic (eds.): *The Social Amplification of Risk* (Cambridge University Press: Cambridge, UK), pp. 13-46.
- Knight, F. H. 1921. *Risk, Uncertainty and Profit*. Washington DC: BeardBooks, Reprinted 2002.
- Lowrance, W. 1976. *Of Acceptable Risk – Science and the Determination of Safety*. Los Altos: William Kaufmann Inc.
- Lindley, D. 1985. *Making Decisions*. N.Y.: Wiley.
- Levy, H. and Sarnat M. 1990. *Capital Investment and Financial Decisions*. Fourth edition. N.Y.: Prentice Hall
- Löfstedt, R. E. 2003. The Precautionary Principle: Risk, Regulation and Politics. *Trans IchemE*, 81, 36-43.
- PSA 2001. Risk Management Regulations. Petroleum Safety Authority Norway.

-
- Renn, O. 1992. Concepts of Risk: A Classification, in: S. Krimsky and D. Golding (eds.): *Social Theories of Risk*. (Praeger: Westport), pp. 53-79.
- Renn, O. 2005. Risk Governance: Towards an Integrative Approach. White Paper No. 1, written by Ortwin Renn with an Annex by Peter Graham (International Risk Governance Council: Geneva 2005).
- Renn, O. 2008. *Risk Governance*. Earthscan. London 2008.
- Rosa, E.A. 1998. Metatheoretical Foundations for Post-Normal Risk, *J. of Risk Research*, 1: 15-44.
- Rosa, E.A. 2003. The Logical Structure of the Social Amplification of Risk Framework (SARF); Metatheoretical Foundations and Policy Implications, in: N. Pidgeon; R.E. Kasperson and P. Slovic (eds.): *The Social Amplification of Risk*. (Cambridge University Press: Cambridge, UK), pp. 47-79.
- Sandin, P. 1999. Dimensions of the precautionary principle. *Human and Ecological Risk Assessment*, 5: 889-907.
- Sandøy, M., Aven, T. and Ford, D. 2005. On integrating risk perspectives in project management. *Risk Management: an International Journal*, 7: 7-21.
- Slovic, P. 1992. Perception of risk: Reflections on the psychometric paradigm. In S. Krimsky & D. Golding (Eds.), *Social theories of risk* (pp. 117-152). New York: Praeger.
- Taleb, N. N. 2007. *The Black Swan: The Impact of the Highly Improbable*, London: Penguin.
- Tversky, A., and Kahneman, D. 1992. advances in prospect theory: cumulative representation of uncertainty. *J. of Risk and Uncertainty*. 5: 297-323.
- Willis, H.H. 2007. Guiding resource allocations based on terrorism Risk. *Risk Analysis*, 27: 597-606.