Jostein Lillestøl,
NHH: 2012/2016

RISK MANAGEMENT AND SAFETY

- An introduction

“Risk"

To laugh, is to risk playing the fool
To weep, is to risk appearing sentimental
To reach out for another, is to risk involvement
To expose feelings, is to risk exposing our true selves
To put your ideas, your dreams, before the crowd is to risk loss
To love, is to risk not being loved in return
To live, is to risk dying
To hope, is to risk despair
To try at all, is to risk failure But risk must be taken
Because the greatest hazard in life is to risk nothing
The person who risks nothing, does nothing, has nothing, is nothing
They may avoid suffering and sorrow, but they simply cannot learn, change, feel, grow, love, live...
Chained by their attitudes they are slaves
Only the person who risks is free!

(Hugh Prather)

Contents

1 Concepts, framework and overview

This chapter tries to explain what risk management is all about, providing framework and concepts.

1.1 What is risk management?

Risk management in some sense is part of most human activities, often more or less unconscious and without giving it a name. For those who say they are doing risk management in their job, there may be huge differences between them, both conceptually and in the task they face and the methods they use. This depends largely on the actual type of business and the context within the business. There may also be differences as to what degree risk management is vital, encompassing and systematic, and whether the activity is there to fulfil some regulatory requirement.

Risk and opportunity go hand in hand, and most often an individual, an enterprise or a nation cannot achieve anything without taking some risks: "Risk in itself is not bad; risk is essential to progress, and failure is sometimes a key part of learning. But we must learn to balance the possible negative consequences of risk against the potential benefits of its associated opportunity" (Van Scoy, 1992).

Two historically important contexts for risk management are:

- Project/industrial risk management.
- Business/finance risk management.

Risk management requires risk analysis. Within each context there are theories and methods for risk analysis, with different origin and developed largely separately by engineers and economists. Concepts, ideas and methods from probability and statistics have to some extent contributed to both areas. There is a lot of common ground in the developments, and in later years we see more tendencies to learn from each other. While earlier theories and methods focused mainly on the negative side of risk, the emphasis is now more on the balance between risk and opportunity.

We may also find risk analysis in other specific contexts, for instance in insurance when judging and pricing different types of contracts, and in medicine when choosing between treatment methods (survival, side effects etc). These are fields requiring a good analytical expertise, offered by actuaries and biostatisticians respectively. They also share some ground with common risk management theories. Again probability calculus and (mathematical) statistics may be put to use. Other fields of potential application are on the national level in services like transport, utilities and public services. On the international level, we have the handling of emissions and other environmental risks. Typical questions asked, in general, are:

- What are the risks (and opportunities)?
- Is it possible to manage the uncovered risks?
- How to describe and communicate these risks?
- How to describe the uncertainties?
- How to weigh the uncertainties?
- How to determine acceptable risk?

A good balanced introduction to risk management in the industrial context, with some side views to business and finance are given by Aven (2002) and Aven & Vinnem (2007)[1]

The risks facing a business enterprise may be of many kinds, among them:

- Strategic risk, financial risk, market risk, operational risk, business continuity and recovery risk, product risk, technical risk, marketing risk, project risk, human safety risk, legal and contract risk, loss of reputation risk, fraud risk, IT risk, counter-spy risk, terrorism risk.

Of course, most risks studied from the operational viewpoint, like they do in an industrial/project setting, may affect the bottom line. Some have traditionally been handled by other than business managers, even if they are key issues in business decisions. They may range from the risk of projects not being finished in time to pollution risks. Until recently, business managers may have thought of risk management as merely a monetary matter. However, the management have to weigh non-monetary issues with economics, and they also have the responsibility to create an environment where this is likely to happen. For people trained in economics, facing other risks than the ones they have learned to state in monetary terms, the questions to be asked may be:

– Do our models take non-monetary risks into account?
– Is it possible to bring such risks into focus, and deal with it rationally?
– How should we balance these risks and economy?
– Can tools like cost-benefit analysis, utility theory and multi-criteria decision theory help?

To be successful, risk management needs to be handled like another management process and be given its place in the strategy of the company, with the full attention of top management. Key operational indicators (metrics) should be used to track and improve performance by managing the aspects of risk that affect employees, associates, customers and shareholders. In recent years the term enterprise risk management (ERM) has emerged, and many organizations have incorporated ERM into a new governance paradigm, in which risk exposure is better understood and managed. They may even have a chief risk officer (CRO) responsible for the whole ERM process of the company, having separate processes for each risk category. Broad categories common to many are: Market risk, operational risk and financial risk.

Risk management has also come to the forefront in the public sector, e.g. in health care and in transportation. Municipalities, counties and national authorities make regulations involving risk, approve and control risk activities and act when serious adverse events to individuals or the public occur. Some of the risk types listed above for private enterprises are also relevant in the public sector, but here more emphasis is on health, environment and safety, and societal risks.

We cannot deal with all of this in these lectures, but will limit ourselves to

- Risk management and safety in general: Concepts, framework and overview (Part 1)
- Approaches and tool for risk management (Part 2)
- Special topics and cases from specific areas (Part 3 and 4)

[1] Aven: Foundations of Risk Analysis, Wiley 2002.
Aven & Vinnem: Risk management with Applications from the Offshore Petroleum Industry, Springer 2007.

1.2 Some risk terminology

Risk

Different fields may have adopted different definitions. This one captures fairly well what we have in mind in general:

Definition: The risk of an activity is the combination of possible consequences and associated uncertainties i.e.

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

where C = Consequences of the activity, U = Uncertainties about C.

This definition is not limited to negative consequences, but encompass potential creation of value by risk taking. Risk management is then to balance between creating value and preventing setbacks.

Remarks. Be aware that there may be differences in the choice of words. Some use Outcome instead of Consequences. However, this may give the impression of just the final result, while all that happens in the chain leading to this is left out. Some use Exposure instead, since you may be exposed to a risk without knowing it, and maybe never get to know that you have been.

A possible definition that widens the scope further is:

$$\text{Risk} = (B, U) + (C_B, U)$$

where B = Possible incidence or initiating events, U = Uncertainty and
C_B = Possible consequences, given initiating events. Here we may name the second sum-term Vulnerability, in particular when we have mostly negative consequences in mind.

There is a difference between how engineers and economists have used the notion risk in the past. Engineers have typically imagined risk as consequence multiplied by probability, i.e. related to expected value, while economists typically image risk as the departure from expected value. Note also that economists, in some contexts, have used the notion risk in situations where probabilities are known (or estimated) and uncertainty when probabilities ("state of the world") are unknown, in order to distinguish the two situations. These notions of risk are too limited to provide a common useful framework for enterprise risk management. How to quantify and interpret risk and uncertainty is a question of choice of a useful paradigm, and we will return to that in the next section.

Risk management

A possible definition of risk management is:

– The systematic application of managerial policies, procedures and practices to the task of analysing, evaluating, controlling and communicating about risk issues.

Here is a formulation of a nationally preferred strategy and approach to risk issues, the Smart regulation - A regulatory strategy for Canada (2004):

> "Risk management is a systematic approach to set the best course of action under uncertainty by identifying, understanding, assessing, prioritizing, acting on and communication about potential threats, whether they affect the public's social, financial or economic well-being, health and safety or the environment".

Risk management is, like most management processes, characterized by steps like:

1. Describe the situation (formulate the problem)
2. Determine goals
3. Seek (alternative) solutions
4. Analysis and judgement of consequences
5. Choice of solution
6. Realization
7. Evaluation

ISO terminology

The terminology used in risk contexts has differed considerably among fields and professions, and have often led to misunderstanding (and added risk). In order to avoid this, the International Standards Organization (ISO) has provided a guide on terminology: ISO Guide 73: 2009 Risk management – Vocabulary (an update of the 2002 version). Here about 40 terms related to risk are defined. This is helpful to prevent confusion among the many stakeholders affected by risk. The terms are the basis for the development of a general risk management standard, as well as being input to standards for specific areas, under way or revision.

The general ISO risk management standard named "ISO 31000: 2009 Risk management – Principles and guidelines" existed as first draft in 2005 and was planned voted on and finalized by 2009. The three main sections of the standard are: Principles for managing risk (clause 3), framework for managing risk (clause 4) and the process of managing risk (clause 5).

The standard states 11 principles for managing risk (clause 3). Risk management should:

1. create value
2. be an integral part of the organizational process
3. be part of decision making
4. explicitly address uncertainty
5. be systematic and structured
6. be based on the best possible information
7. be tailored to the context
8. take into account human factors
9. be transparent and inclusive
10. be dynamic, iterative and responsive to change
11. be capable of continual improvement and enhancement

Risk management should be an integral part of the organization supported by management. The standard advocates a framework for managing risk (clause 4) by means of a risk management process (clause 5), to be used at different levels and in different contexts. This framework should ensure that risk information is derived from these processes, and is adequately reported and used for decision making at the relevant organizational levels. The clause also gives guidelines for designing, implementing, monitoring such a management framework

The following exhibit illustrates the components of the framework and its connection to the risk management process:

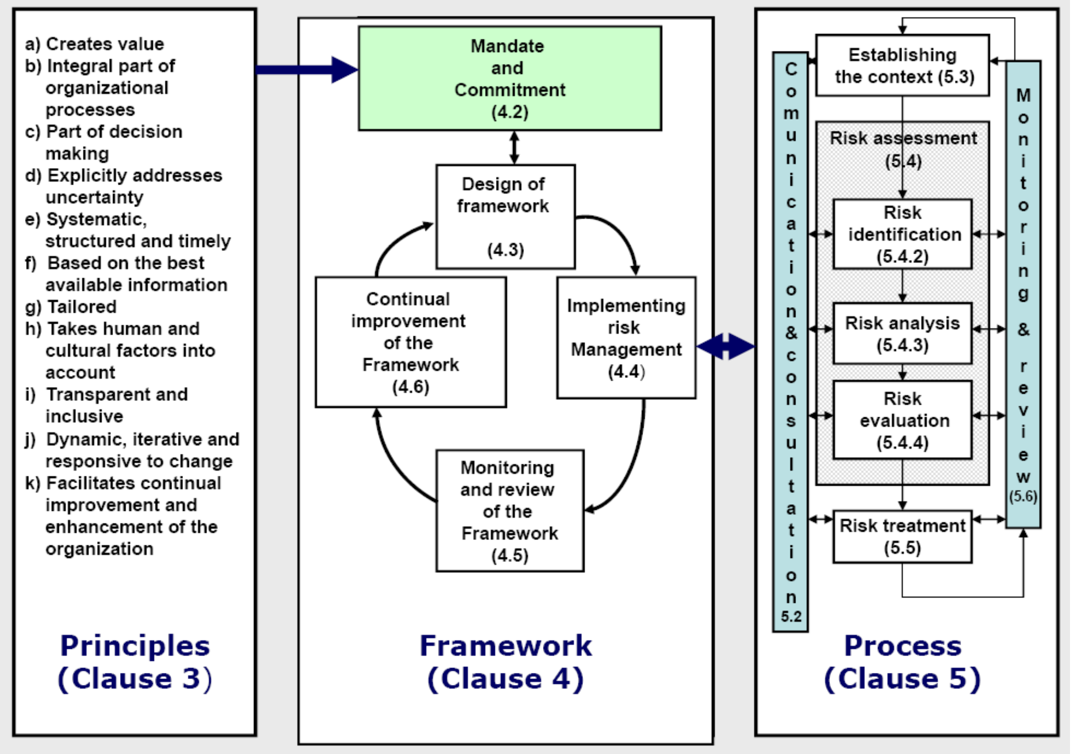

ISO 31000:2009 Figure 1 – Relationship between the principles, framework and process

Concerning the risk management process, the terminology shall be understood as follows: Risk assessment is the combination of risk identification, risk analysis and risk evaluation, where risk identification and analysis is the systematic analytical work undertaken and risk evaluation is the key decision-making steps based on the analysis. Risk treatment is the management of acceptable risk. We will return to these activities in section 2.1.

Exercise
Try to find out in some detail what ISO 31000 says about risk assessment.

There are many ISO other standards with strong emphasis on risks, among others the ISO 9000 series on quality management, the ISO 14000 series on environmental management, and the ISO 27000 series on information security management. Moreover there are standards for specific industries dealing with their specific risks, e.g. food, construction, chemicals etc.

1.3 Uncertainty and probability: Choice of paradigm

As stated above, risk is the combination of uncertainty and consequences. Uncertainties can be expressed by probabilities. A fundamental issue is then whether risk should be viewed as an objective entity, something inherent in the situation, there to be uncovered if we have sufficient data and appropriate analytic skills. The issue is both philosophical and practical, and affects how we should go about to analyze and manage risk, and how we should interpret, accept or challenge a risk analysis presented to us.

For the probability P(A) of an outcome A we have mainly two different interpretations:

Interpretation of P(A)	Implications
(a) Long term fraction of outcome A in independent repeats of opportunity to observe (i.e. being independent of observer).	The underlying P(A) is taken as unknown to be estimated by (limited) data (with the help of some statistical model). Pretends to be objective*
(b) Measure expressing the uncertainty of an analyst about A to happen, based on some background information and knowledge.	There is no true probability, the probability may depend on the analyst(s) Appear to be (too) subjective

* In some cases, these probabilities may be arrived at by symmetries or by design without observation e.g. coin, dice and card games and in lotteries. However, this is outside our scope.

Classical risk analysis has conceptually stayed close to (a), but there are some problems with this.

<u>Problems with (a)</u>:

- Most risk analyses are performed in situations that are not fully repeatable.
- The objectivity may be illusory, since model assumptions have to be made
- We often have scarce data, so that the “true” probability is estimated with uncertainty
- Although classical statistical theory provides error limits, this adds a new layer of uncertainty
- Give room for “experts”, hard to challenge, since they are “objective”

These problems have led many to leave (a) as paradigm for risk analysis and adopt (b) instead, among them Aven op.cit. This means that the following is appreciated:

<u>Advantages with (b):</u>

- Does not give false impressions of objectivity
- Gets rid of the extra layer of uncertainty
- May encompass objective reasoning when hard data exist
- May more easily take perceived risks into account
- Risk analysts may be more easily challenged

Taking (b) as paradigm we have implicitly accepted that there is no true risk, but the assigned risk may depend on the reporter(s). On the other hand, the risk experts are now taken down from the ivory tower. Most risk analyses of significant importance to many stakeholders have to be performed by a group of people with diverse and relevant insight and/or competence. For the rest of us, it is a matter of trust. The major drawback may be that this leaves the field more open for anyone to pour out unfounded doubt to anything that goes against their interests.

Again, risk is the combination of uncertainty and consequence. We then face the question of expressing the risk in relevant and meaningful terms. For a consequence measured numerically having different outcomes that are assigned probabilities, we may compute the expectation, i.e. the sum of the outcomes weighed by their probabilities. Within the classic paradigm this can be interpreted as a "long run average outcome" in repeats. This may be relevant in some contexts, but definitely not in others. Within the alternative paradigm it becomes a judgment, or derived from judgments, given some background information, which may be specific to the analyst(s). In some contexts, for example some contexts in finance, the deviation from expectations, often named volatility, is the key risk quantity. In cases of hazards, the extreme of the probability distribution is of primary interest. In fact, this is one of our prime concerns in these lectures. In any case, there is a vital conceptual difference between the interpretations of any expression of risk within the two paradigms. In particular we see the limits of the classical one for extreme outcomes with low probability, like oil spills and nuclear accidents. They occur so infrequent that thinking in terms of repeats under the same conditions is not very helpful at all.

The choice of paradigm also has implications for how we interpret and judge historic data to throw light on the future. Within the classic paradigm with true underlying risk parameters, the question becomes how to estimate them with confidence. Within the alternative paradigm the question is rather one of best possible prediction of observable future quantities. In the first case we are close to statistical theory as taught in most introductory statistics courses of today (estimation and confidence intervals etc.). In the second case we are close to predictive ideas, closer to Bayesian statistics. In this case it has no meaning to speak about the uncertainty in the assigned risk probabilities, they are given based on some background information, period!

Opposing perspectives on risk that bears some relation to this may be found in social science literature. They are mainly

The techno-scientific perspective ("risk-realism"):
— Risk exist as real world, regardless of beliefs,
— to be uncovered by analytic (scientific) means,
— by researchers and professionals,
— with other personal views serving as additional input.
— Risk is measurable, calculable, manipulable and profitable

The socio-cultural perspective ("risk-constructivism")
— Risks are mental constructs,
— never fully objective or knowable outside of belief systems,
— embedded in our culture, which give them meaning,
— the political, moral, and aesthetic judgments produce risks
— Risk must reflect the public concern.

Opposing statements like this "Realism ignores the system level, and is at best naïve!" and "Constructivism is absurd, and may increase risks, due to misguided perceptions in the media and the public!" Well, constructivism may provide some needed social critique, but does not provide any useful framework for analyzing specific risks.
These notes will take a middle ground, but stay close to the realist view, if possible and reasonable, and then point to possibilities at the system level.

Exercise: Reflect on the paradigm in the probability and statistics courses you have taken.

Let us initiate some additional reflections: A dice is going to be rolled and the question is

Q1: Will the dice show five?

Suppose the dice is rolled and the outcome is not revealed to you, but written down on a piece of paper. Then the question is:

Q2: Is the number written down a five?

Most people are comfortable to answer Q1 by something like "one in six" either, having in mind, a long run frequency interpretation or a symmetry argument. Some may have problems with the same answer to Q2, arguing that now the outcome is there as a fact, and unrevealed to us it is either 100% true or 0% true (i.e. false). The difference between the two situations is that the first represents genuine uncertainty about an event going to happen (randomness), while for the second the event has taken place, but we are kept in the dark about the outcome. However, there are compelling reasons for treating the two situations the same way. Moreover, since the information available may differ between people in practical situations (of more relevance than this), it seems natural to regarded stated probabilities as personal. We may still use a symmetry argument treating the six numbers as equally likely (the principle of insufficient reason). An example from the courtroom: The defendant is either innocent or guilty. The act has happened, and guilt is either 0% or 100%. Expert witnesses and jurors may express (or have in their minds) in between probabilities depending on their knowledge (or ability to grasp the evidence).

We will assume that the reader has a basic knowledge of probability calculus, including the addition rule, multiplication rule for independent events, conditional probability and Bayes law. To recall some of this please try the next exercises.

Exercise: Bayes law
On a given night the chance is one in ten thousand that a burglar will enter your house. In that case it is 95% chance that the installed alarm will turn on. However, there is also a 1% chance that the alarm will turn on for no apparent reason on any given night. During one night the alarm turned on. What is the probability that there is a burglar in the house? Repeat the exercise for an alarm less sensitive to other circumstance, replacing the 1% by 0.1%. Comment on the results.

Exercise: Bayes updating
Remaining lifetime Y is heavily dependent on the patient's health state θ and may be modelled conditional on the state: Assume three states: $\theta = 2$ (seriously ill) $\theta = 1$ (moderately ill) and $\theta = 0$ (not ill) with prior probabilities reflecting the state of incoming patients

$$P(\theta = 0) = 0.85 \quad P(\theta = 1) = 0.10 \quad P(\theta = 2) = 0.05$$

(a) Suppose a test for indicating serious illness (+) or not (–) is available, and from experience the probability of indication is dependent on the patient state as follows

$$P(+ \mid \theta = 0) = 0.10 \quad P(+ \mid \theta = 1) = 0.50 \quad P(+ \mid \theta = 2) = 0.90$$

Use Bayes law to show that the posterior state probabilities are

$$P(\theta = 0 \mid +) = 0.472 \quad P(\theta = 1 \mid +) = 0.278 \quad P(\theta = 2 \mid +) = 0.250$$

(b) Assume that the test is repeated with results independent of each other
Calculate $P(\theta = 2 \mid ++)$ by two different methods
(i) By one more step from the posterior
(ii) By one joint step from the prior
Comment on the result.

(c) Suppose that the doctor, based on other evidence at entry, have assigned the prior to

$$P(\theta = 0) = 0.20 \quad P(\theta = 1) = 0.40 \quad P(\theta = 2) = 0.40$$

Redo the calculations in (a) Answer (0.034, 0.345, 0.621)

Exercise: Unfounded practice?
A suburban bus station has a parking lot reserved for park and ride passengers to town. The parking is free, but you are allowed to stay maximum 24 hours. This is controlled by setting a crayon mark on one tire of each car at the point where the tire touches the ground. Every day the cars are checked for marks. If a mark is at the ground spot, it is taken as if the car has not been moved, and the car owner is fined by mail. You parked there each of two consecutive mornings, went to work and returned home. Nevertheless, you got a fine in the mail.
How would you speak up against this practice using probabilistic arguments?

1.4 Human hazards: Some principles

Many business decisions involve some human hazards and even potential fatalities, and plans have to be made to prevent or reduce the risks. The question may be how much to invest in safety. Since economic analysis is the basis for most business decisions, efforts have been made to set a price to human life in order to make it an integral part of the analysis. This is neither easy nor appealing, although it is often done implicitly. In some cases there is a way out along these lines: Consider a big investment project, like the ones in offshore oil business, and an alternative solution is assessed with some added safety. However, this solution is very expensive, and suppose the assessment tells that the present value of the cost per averted (statistical) life lost is $500 mill. If the calculations are trusted, this might be sufficient for deciding to drop it, and this is accepted by all stakeholders.

The Cautionary principle

In many cases one may advocate that a cautionary principle should prevail for issues related to human hazards and fatalities, even if assigned probabilities are small and economic cost-benefit analysis does not support additional measures. Cautionary means going beyond the calculatory cost-benefits in implementing such things as:

1. Extra component quality, redundancy and back-up
2. Robust design of components or systems and “fool-proofing”
3. Extended maintenance routines
4. Supervision, e.g. by suitable alarms
5. Safety barriers, i.e. prevent an initiating adverse event to develop
6. Training of personnel

In some cases, the cautionary principle may lead to abandoning the activity.

The ALARP-principle

In cases where an activity involving human hazards is required, but the question is how much caution should be invested, there is a need for requiring a reduction of the risk to a level “As Low As Reasonably Practicable”. The ALARP-principle looks rather vague at first sight, but it is found very useful in practice. More on this later.

The Precautionary principle

When we, to large extent, lack knowledge of the consequences of an activity, it should not be carried out. We then follow a precautionary principle. There are different opinions on how far the principle extends, and various definitions of the principle exist. United Nations Educational, Scientific and Cultural Organization (UNESCO) published in 2005 “The Precautionary Principle”, a document prepared by the World Commission on the Ethics of Scientific Knowledge and Technology (COMEST). Here the principle is stated as follows:

“When human activities may lead to morally unacceptable harm that is scientifically plausible but uncertain, actions shall be taken to avoid or diminish that harm”.

Here it may not be clear what "uncertain" means. Some alternatives are: (i) unknown type of consequences, or (ii) known consequences, but unable to predict the extent, or (iii) unable to establish the dependence of the consequences on underlying factors. As indicated in the definition, this becomes a question of values and ethics. The precautionary principle has, in particular, a role to play in questions related to long-term risks, for example related to public health, and to environmental issues. In this context the precautionary principle applies when

- There is considerable lack of scientific knowledge, i.e. about causality, magnitude, probability and nature of harm.
- The possibility of unacceptable harm is established: Brought forward scientifically, but not easily refuted scientifically.

Situations when harmful consequences are understood, but known to be not likely, belong to regular risk-based analysis. Appealing to the precautionary principle is then a misuse that deprive the concept its specific role

Acceptable risk?

The issue of acceptable risk can be discussed in different contexts, among them

(i) Assessment of risk in view of preset maximum risk levels in the given application area
(ii) Comparison of assessed risk with accepted risks in other application areas

We consider here the first context. Although such levels may be set with the best intentions, they may have an adverse effect. It may shift the focus from developing a broad understanding of the risk picture to just fulfilling the requirement. No emphasis on risk reduction beyond the preset limits, may work against continuous improvement efforts. The ALARP principle is, in a sense, a middle way between the practice of preset limits and continuous improvement ideology.

Norway has been a safety pioneer in the offshore oil and gas sector since the 1970's, with legislature based on detailed prescriptive requirements. An example is the so-called 10^{-4} criterion from 1980, which is the maximum probability per year for each of nine types of accidents. This initiated the early systematic risk analysis efforts, with positive effects, but which in some cases degenerated to a numbers game. Norwegian authorities have gradually realized the possible adverse effects of preset criteria, and have moved towards ALARP ideas in their legislature, more like the UK legislature. Now the regulations have more emphasis on the risk analysis process, where the operators are made responsible for defining safety objectives and risk acceptance criteria pertaining major accidents and hazards to the environment.

<u>Example 1:</u>
Probability that an individual get killed in accident during one year < 0.01%

<u>Example 2</u>:
Fatal Accident Rate FAR < 10 for all personnel on the installation
(FAR= Expected no. of fatalities per 100 mill. exposure hours)

However, skeptics say that the change of legislature has not led Norwegian operators to adopt a risk reduction mindset, and that it is still an effort just to fulfill some criterion, and if major improvement possibilities are revealed, they are often dismissed by cost-effectiveness arguments. On the other hand, they question supervision authorities who are not likely to use improvement efforts as a separate criterion and check for that.

The ALARP-principle: Implementation

For a given risk measure R we may define two risk levels $R_1 < R_2$, representing a negligible and intolerable risk so that

1. If $R < R_1$ the risk is taken as negligible
2. If $R > R_2$ the risk is taken as intolerable
3. If $R_1 < R < R_2$ apply the ALARP-principle: "as low as reasonable practicable"

The latter is the most frequent situation in practice, and then one proceeds to search for risk reducing opportunities, preferably by means of some well-defined procedure. The gap $R-R_1$ may indicate how much effort should be spent (perhaps in relation to R_2-R_1). The ALARP-principle then means that one should implement a risk reduction opportunity, unless it is documented that the costs are unreasonable high, compared to the statistical life saved. The question is then: What is high? In some countries such numbers may be stated, and may differ between activity areas. In the UK costs similar to NOK 25 mill. are stated for the transport sector, while NOK 75 mill. for the offshore oil sector. In Norway it is known that investments above NOK 200 mill. per saved statistical life occur, while societal investments elsewhere are rejected for far less than this, perhaps down to NOK 1 mill. in some areas without being explicit or aware of it.

Formal ALARP-procedures are available and described in the literature. Of course, those who still favor acceptance criteria voice their critics, but strong defense exists (see Aven, 2007). We will only mention that the use of acceptance criteria allows delegation to lower levels in the organization, while ALARP requires a higher level, being able to handle the trade-offs between economics and hazards. Not all decision makers are happy with this challenge.

For human hazards, it may be strikingly differences between the perceived ("subjective") risk and the statistical ("objective") risk. In particular, this is so when the consequences may be dramatic, even if the probabilities are very small. When people have a sharper focus on the consequences than the probabilities, their perceived ranking of various risks will quite often differ wildly from the ranking derived from statistics. To gain acceptance for a decision, it is therefore often more important to limit the consequences than reduce the probabilities.

Exercises

1. Find formulations of the precautionary principle in relation to harm to public health and to the environment.
2. What do European Union law and regulations say about the precautionary principle? Express some critical views.

Many hazards related to health and safety may lead to injuries or other harm to the individuals, without leading to any fatality. There may be a need to define risk metrics that represents such hazards, and preferably allow comparisons with fatality risks. This may be a challenge as

— definitions of "major injury" and "minor injury" are needed, and
— different types of injuries may have to be put on a common scale.

The concept of "Equivalent fatalities" is an effort to make the comparison with fatalities: 1 fatality = 10 major injuries = 100 minor injuries. This choice may of course be questioned, and some feel = 200 minor injuries is more appropriate. The use of such metrics, is to help making some comparisons, and should be spelled out in risk reports, so that they can be challenged by the stakeholders, if they are felt to be inappropriate. In some cases the harm to the individuals at a workplace may be linked to doses, e.g. of toxic gas, heat or overpressure. The concept of "Dangerous dose" is a dose that has the potential of leading to a fatality, but does not necessarily do so. This may be defined as a dose that gives any of the following

— severe distress to everyone
— a substantial number needs medical attention
— some need prolonged medical treatment
— fatal for any highly susceptible

The British Health & Safety Executive (HSE) have suggested this concept, and equated "dangerous dose" with the dose that would kill 1% of a "typical" population exposed to it. Note that this context has nothing to do with risks levels connected to consumer products, e.g. additives. Above we are talking about doses at a workplace at an event to be prevented, and the concept is there to measure and compare risk levels in the risk assessment process.

Long-term human hazards related to the workplace and public health in general is an important area of concern. Here expertise in epidemiology may come to help. We will not go into this in any detail here, but just mention two concepts:

An individual may die (D) from a number of "competing" causes, and it is of interest to have some measure of the risk attributable to each cause. Denote exposure to a specific cause by E, and non-exposure by its complement E^c (sometimes written $\bar{E}$). Two such measures in terms of probabilities are

Relative risk: $$RR = \frac{P(D|E)}{P(D|E^c)}$$

Attributable fraction: $$AF = \frac{P(D) - P(D|E^c)}{P(D)}$$

AF may be interpreted as the fraction of those dying in excess of those who would have died anyway without exposure of E. In other words: The fraction saved by removing the exposure.

Exercises

(a) Show that $$AF = \frac{(RR-1)P(E)}{(RR-1)P(E)+1}$$

(b) Compute AF for RR=2 and each of the cases P(E)=0.05 and P(E)=0.5

We close this section by pointing to some other aspects to human hazards:

In a modern affluent society new risks have arisen, some related to the generation of wealth itself and some due to new life styles. To some extent this has become a concern and challenge for society itself.   We also see that the public is expecting more from the authorities, in particular if taxes are high. However, there are more risks in the world than our society is able to copy with at the same time.

Examples:

- People paricipate in hazardous expeditions or risk sports, and expect to be saved when in troble, even if this endangers other peoples life.
- People on vacation to far away places, and are hit by some natural disaters and expect immediate help from the foreign office.
- Peopl experiencing trauma after incidence, and claim compensation from authorities, even if the society is not at fault, and they luckily escaped the fatality themselves
- People in search of circumstances and a diagnosis that can possibly lead to some benefits from society (often with the help of lawyers)

Some say that we have made the state into an insurance company.
Research findings related to human risks often become headlines in the media. The way the journalist chooses to present the results makes a great difference on how the findings are interpreted by the public. Suppose the headline says the risk of dying within a given time span is doubled if you are exposed to an environmental agent. This sounds rather scaring. However, the reported doubling is typically the relative risk, which should not be given without reference to the absolute risk. In this case taken to be the probability of dying without being exposed, i.e. the denominator in RR. If this is small, it may be nothing to worry about. The double of almost nothing is nothing!   However, this does not create headlines. We are now touching upon the issue of risk literacy, and we will return to this later.

Exercises:

1. A downtown area of a city (Bergen) have frequently bad air, due to releases from heavy traffic and outdated heating systems, in particular on cold winter days when atmospheric conditions are unfavorable and have cause inversion. It is believed that this may increase health problems like asthma. However, in a survey of the 1st grade children in the area (Årstad) the frequency of asthma was lower than other areas of town, including a mostly rural area (Arna). Does this rule out that the bad air may cause increased risk of asthma?
2. You are at the playground with your 2 year old child, and she wants to try the big spiraling slide for the first time. You are very protective, and go sliding with her on your lap. Is this really a protective act that reduces the risk of injury? What kind of data is needed to settle this issue? Answer: See the Well column in New York Times April 23 2012.
3. It is National Day May 17th and you are in the school attending the celebration at your local school. The responsibility for the arrangement is always the parents of the third graders, staging games and lotteries, selling goodies and provide the stage for performances of the kids.   Your kid is in second grade and you decide to look around to be prepared for next year. In particular, you to try to see whether there are risks of some kind.   Among others, you pay attention to the following three situations. Do you see any necessary changes or necessary precautions?

A:

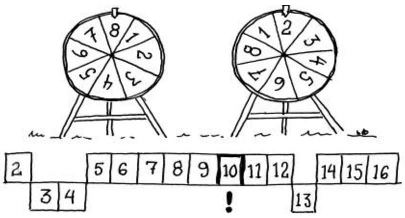

B:

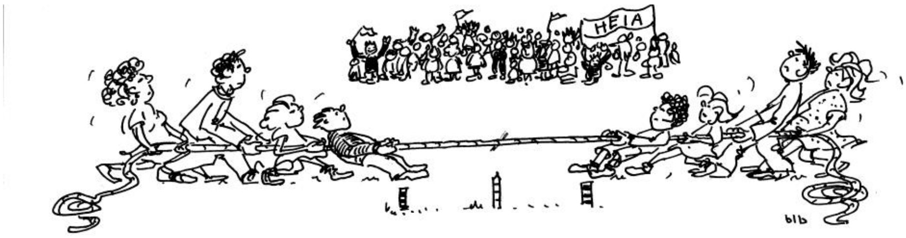

C:

1.5 Ethics and risk assessment

Some examples of unethical risk assessments:

1. A risk assessment is performed to justify a decision already taken.
2. Deliberate failure to expose all possible hazards and/or all possible consequences.
3. Assuming too favorable assumptions, without notice of alternative assumptions leading to a different picture, e.g. assumption of independence.
4. Making use of inappropriate risk criteria.
5. Not doing anything, after an assessment requiring some action.
6. Someone that complies well with the preset risk levels take the opportunity to increase the risks, in order to save time and money, called "reverse ALARP".

An example of 4 may occur in the case of "transient exposure" where many people are exposed a short time. If you divide the risk number of someone being harmed by the large (and sometimes disputable) number of people, you may obtain low risk numbers (called the "salami slicing approach").

Ethics in risk assessment may be looked upon from the perspective of

(i) the purpose or intention of an action
(ii) the consequence of an action

By (i) the question is mostly about what is right or wrong, without reference to costs, while (ii) is about balancing benefits/costs and consequences hard or impossible to measure in economic terms, among them possible hazards and fatalities. In case formal economic analysis cannot help, the ALARP-principle may come to rescue. Furthermore, an action may be given high benefits without hazard risks for some, but high hazard risks for others. This can in some cases be resolved by compensation or insurance, but not always, and quite often both perspectives (i) and (ii) have to be considered concurrently.

We will now look more closely at the ethical side of the common practice of setting acceptance levels for risk. If we are uneasy by using the term acceptable risk, we could replace this by the phrase that we accept a solution involving risk attributes (Aven, 2007).

The setting of acceptance level for risks by regulators may seem natural for (i). Some comments to this: One argument is that if operators are given the opportunity set their own criteria, as in some ALARP-regimes, they will be set so that they are always met, with no pressure to improve. It may appear that (i) with the use of preset acceptance levels requires the idea of true risk, i.e. the classical paradigm. However, even within a subjective Bayesian framework such levels may be used for comparing different subjective risk assignments, and its justification for ethical reasons remains much the same.

The regulator task of setting levels for practical use, is a delicate one, and has to accommodate any paradigm of risk. Suppose that 0.1% probability is regarded as a maximum acceptable true risk for some individual risk. When estimated with uncertainty, the point estimate of the risk has to be much lower in order to guarantee that this is fulfilled (say upper confidence limit less than 0.1%). Should the regulator therefore announce a much lower acceptable risk level, say 0.01%, so that when the assessment is below this, it is guaranteed with high probability to be below the wanted 0.1%? Or, should they just announce the weaker 0.1% limit? Within the Bayesian paradigm where the assessments are subjective and may vary, it is consensus that may provide us with "guarantees", and the issue is not much different.

Most people are willing to take higher risks if they are voluntary than non-voluntary, and in fact of quite different magnitudes. A comparison of the fatality risk of some voluntary activities with the corresponding non-voluntary fatality risk of an industrial worker, may give multiplicative factors like

Industrial worker	1
Car driving	10
Motor biking	100
Mountain climbing	1000

It also happens that for two risks of similar magnitude statistically, one is accepted the other not. Typically, one may be willing to accept things that you are used to. In some areas, the work force is paid extra to take health risks and fatality risks. This may be the case for handling dangerous material or working at dangerous locations, e.g. in a mine, at sea (war zone or deep diving). But quite often this does not happen, even if the risks are well known. In many cases, there are uncertainties about the probabilities and/or the consequences of an activity, and in some, regrettably, partly known, but not revealed and compensated at the time of employment. After adverse events have occurred, there may be some opportunities to compensate, but quite often workers may face unwillingness to accept the causes or to compensate. In fact, no money are likely to compensate fatality or extreme loss of life quality. This has a clear ethical dimension.

The choice of appropriate action in case of possible human hazards will quite often have ethical aspects. Moral philosophers have dealt with this issue over the years, with a variety of different views. There are two main lines of thought: *Consequentialism* and *Deontologism* (deon="duty" in greek). Consequentialism means that moral assessment of acts or intentions are based only on their consequences, while deontologism is normative theories on acts that are morally forbidden, permitted or required. Here one accepts that there are some acts that cannot be justified, no matter how morally well the consequences turn out to be. Some deontoloists take the agent view (i.e. a desire to keep own house in order) and some the patient/victim view (i.e. they have rights). Another dimension of deontologism is a felt need to have a well described social contract, based on principles that most people would accept, and from which acts may be judged as morally justified or unjustified, so called *Contractualism*. Many issues related to risk can be discussed within the different paradigms. Many examples in the literature are, although interesting, more an exploration how far the ideas may extend. Example: Kill one person, remove vital organs and transplant to in 10 people who would otherwise die shortly after.

Discussion: There are three different groups with risk of dying 100% 50% and 10%. A drug is available which lower the risks by 10% to 90%, 40% and 0% respectively. However, the drug is available in limited quantities, just enough for only one group. Who should be given the drug?

1.6 Health, Environment and Safety (HES) at the workplace

Enterprises in Norway have to relate to the following:

- The Working Environment Act (“Arbeidsmiljøloven”)
- The Internal Control Regulations (“Internkontrollforskriften”)

The Working Environment Act (of June 17, 2005 with later amendments) relates to employment and working conditions. A prime purpose of the Act is ”to secure a working environment that provides a basis for healthy and meaningful working situation, that affords full safety from physical and mental influences and that has a standard of welfare at all times consistent with the level of technological and social development of society”. The Working Environment Act has 20 chapter and many of them relates to health, environment and safety (HES), in Norwegian: ”Helse miljø og sikkerhet” (HMS).

The Internal Control Regulations (effective of January 1997) is about the implementation in public and private enterprises of The Working Environment Act and other legislation that relates to the working environment as well as the consequences of the activities of the enterprise for customers (products and services) and for the society (e.g. The Pollution Control Act). It requires enterprises to have systematic approach to health, environment and safety issues, and be able to document this to the stakeholders, among them as supervisory body Norwegian Labour Inspection Authority (“Arbeidstilsynet”). Webpage: http://www.arbeidstilsynet.no/

Let us look at some details, with emphasis on those related to HES and hazards:

The Working Environment Act spells out in chapter 2 the main duties of the employer and employee with respect to HES work, followed by chapter 3 on working environment measures, stating, among others, that the employer is responsible for systematic HES work by

- establishing goals in the area, having an overall view on responsibilities, tasks and authorities,
- surveying hazards and, on this basis, assess risk factors, prepare plans and implement measures to reduce the risks,
- ensuring continuous control of the identified risk factors, and
- reviewing the HES work, to see that it works as intended.

This also includes occupational health services and employee participation, among others by a specific working environment committee. Chapter 4 starts by stating that the working environment of the undertaking shall be fully satisfactory when the factors that may influence the employees’ physical and mental health and welfare are judged individually and collectively. The standard of safety, health and working environment shall be continuously developed and improved in accordance with developments in society.

In these chapters, also a number of special precautions are detailed with respect to the physical and, the psycho-social environment. Among the physical issues mentioned are: Buildings, equipment, indoor climate, lighting, noise and radiation, chemical and biological hazards.

Mentioned are also the avoidance of physical strain and monotonous repetitive work. Work equipment should be designed and provided with safety devices so that the employees are protected against injuries.

Chapter 5 deals with the obligations to record and notify. Employers have to keep records on work related injuries and diseases which have to be made accessible to the Labor Inspection Authority, safety representatives, occupational health services and the working environment committee. Statistical records should also be kept on absence due to sickness.

The following two chapters give more details on the responsibilities of safety representatives (chapter 6) and the working environment cooperation (chapter 7)

The Internal Control Regulations requires the roles and responsibilities for the health and safety issues in the enterprise to be clarified. Risk analysis and assessment must be carried out, and plans of action made and carried out according to assessments. Those responsible for the enterprise must ensure that internal control is introduced and executed in collaboration with the employees and their representatives. The extent of the HES work will depend on the risk picture as well as the size of the work force, as shown in the graph below. With small workforce and low risks, the HES activities may be straightforward, while large in one or both respects may require extensive and/or more sophisticated HES work.

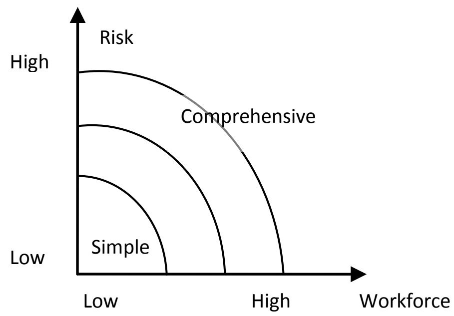

The regulation contains the following

1. Getting started

   - Taking the initiative
   - Providing information and motivation
   - Setting objectives and defining responsibilities and lines of delegation
   - Organizing and planning the introduction

2. Identifying problems/obtaining overviews

   - Obtaining an overview of relevant laws and regulations
   - Taking stock of existing health, environment and safety routines
   - Systematization and safekeeping of documents
   - Identifying health, environment and safety problems

3. Planning and ranking measures

- Drawing up an action plan for implementation

4. Follow up

    - Implementing measures
    - Make improvement work a natural part of operations
    - Rectifying errors and defects
    - Periodical review of health, environment and safety activities

The Labour Inspection Authority offers advice on how to implement this, and provides several schematic tools for identifying hazards and prioritize measures to control the risk.

Exercises

1. Get an overview of the HES material on the web site of the Labour Inspection Authority Alternatively check the web site of the Health and Safety Executive www.hse.gov.uk
2. Examine some details of the Internal Control Regulations, in particular the identification of HES problems
3. Go to the web site of some major Norwegian companies and find out what they have to tell about their HES activities (e.g. Statoil, Telenor, SAS, Statkraft).

There is a Norwegian standard for risk assessment NS-5814: 2008 (”Krav til risikovurdering”) describing the joint process of planning, risk analysis and risk evaluation (cf. the ISO 31000 standard described in section 1.2). NS-5814 describes some requirements for a the elements in this process and also the role of risk assessment in risk management and the factors that may affect the planning and execution of the risk evaluation, among others, establishing risk acceptance criteria. The revision of this standard in 2008 puts stronger emphasis on the planning phase and with the risk evaluation as a new element. The standard is intended for sectors, business or public, which do not have specific standards for risk assessment of their own.

The HES-work will typically involve a description of the situation, of what is being done, of the outcome, and of how one would like it, as depicted in the following diagram:

This is a simple example of a so-called influence diagram (ID), also named relevance diagram. It is an intuitive way to identify and display the essential elements of the situation, including objectives, uncertainties and decisions, and how they influence each other. The idea is to obtain a high-level conceptual view, which may provide the basis for a more detailed description and modelling. ID is suitable for teamwork, since it can handle incomplete sharing of information among team members to be modeled. [2] Here is a simple influence diagram related to risk analysis:

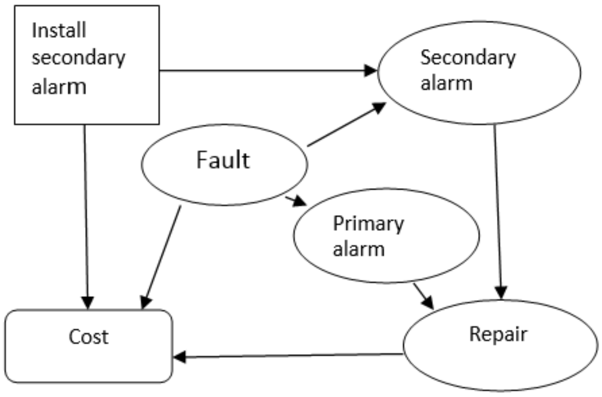

Some of the other tools, graphical and schematic are described in section 2.2. Here we just give an example of a crude schematic tool for dealing with hazards (for both staff and customers) at a restaurant.

Example: HES in a restaurant

What are the hazards?	Who might be harmed and how?	What may reduce the hazard?	Action required			
			priority	**by when**	**by whom**	**done**
Trips and fall – stairs, objects	Staff – ... Guests – ...	Improve lightning Replace carpet ⋮	⋮	⋮	⋮	⋮
Slips and spillage – kitchen – while serving	Staff – ... Guests – ...	Better housekeeping ⋮		⋮	⋮	
⋮						
⋮						

Exercise: Fill in more hazards in the scheme!

[2] Influence diagrams may be seen as an alternative to so-called decision trees, and more suitable in cases when the situation is complicated, and not easily described by a tree-structure. As an influence diagram may include probabilities and decisions, it may be seen as the basis for so-called Bayesian networks (see section 3.11).

The human factor

Although many systems are highly technical, they involve human interaction. Humans do errors and contribute to adverse events, either as a contributing root cause or lack of adequate response to developing events. Quite often, the human factor is vital to the understanding of how a system works. Efforts to take human error probabilities (HEP) into account in project and system risk analysis are faced with a number of problems:

(i) Data are scarce and limited to simple definable situations

(ii) Frequency based data may fit into a context of repeatable tasks, but most human situations are not repeatable. People learn from their mistakes, and people hardly meet the same situation with the same level of knowledge and experience.

It may be useful to think in terms three different performance levels:

1. Skill-based performance
   - Follow a prescribed procedure
2. Risk-based performance
   - React to changing circumstances (if … then)
3. Knowledge-based performance
   - Act in new situations based on knowledge

Some types of errors are:

- Omission: Step in task (or whole task)
- Selection: Wrong choice
- Sequence: Wrong order
- Timing: Too early – Too late
- Quantity: Too little - Too much

Human reaction to critical situations differs wildly, from those who act quickly and intuitively to those who to take time to reflect on the options. Organizations involved in risky operations typically have emergency plans, and employees are expected to follow established and rehearsed procedures. However, situations may occur that are not described in manuals, and if they are, there may be no time to consult the manuals or consult a superior. In military terms: "When the planned operation is launched, the plans are history, and the officers (and soldiers) are likely to face situations not covered by the plans". What kind of human capability and acting style are the most valuable when facing critical situation that requires quick decisions? We may contrast two styles with different characterizing elements:

<u>The analytic style:</u> Reflective, controlled, step by step, detailed, explainable.
<u>The intuitive style</u>: Quick, automatic, "gut feeling", not so explainable.

The research on intuition and decision psychology is divided. Here we briefly contrast two writers: Kahneman and Gigenrenzer. [3]

Kahneman (Nobel laureate) warns against regarding intuition and heuristics as a virtue. In particular, he has researched into how humans react to uncertainties, and found that they may be very inconsistent, often due to biased interpretations of data. This may have serious consequences for decisions based on limited experience.

Gigenrenzer argues that, for complicated situations which require a quick response, heuristics may make the situation less complicated and thus provide a more effective decision than an analytic approach. He points out that in many cases less information is better than more, since the extra information is more likely to be just “noise”.

Effective leadership may require ability to combine both styles, and be conscious when to follow a particular one. Some say that scoring high on analytic abilities and low on intuitive is preferred to scoring low on analytic and high on intuitive, and some say the opposite. However, intuition has to be based on knowledge and experience, and it helps to be trained in taking intuitive decisions. (cf. recent Norwegian research into this area by Bakken 2011)

[3] Kahneman (2011) Thinking Fast and Slow, Gigenrenzer (2007) Gut Feelings: The intelligence of the unconscious

1.7 Some risk statistics

Risk statistics may be obtained from various sources: From Statistics Norway ("Statistisk Sentralbyrå"), from The Norwegian Labour Inspection Authority ("Arbeidstilsynet"), from Petroleum Safety Authority ("Petroleumstilsynet") or from other specific supervisory authorities. Here is a graph from the website of "Arbeidstilsynet" exhibiting yearly work related fatalities in Norway in the period 1964-2007:

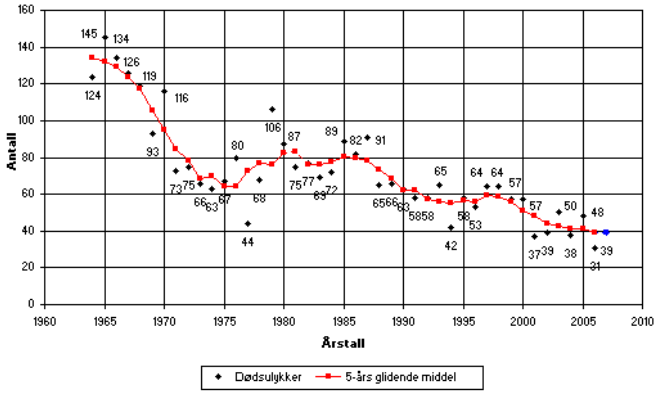

Work related fatalities in Norway 1964-2007

We see that the number of work related fatalities in Norway has declined during the period, rapidly until the mid 1970's, then and standstill until the mid 1980's, and then declining again after that, probably due to greater awareness and introduction of HES. However, some of the decline may be due to fewer people working in hazardous environments. To tell the effect of HES work, we need statistics over time from specific occupations. The standstill in the 1980's could probably be due to the pioneer days of North-Sea oil. However, the offshore oil and gas activities itself were outside the supervision of this authority and reported elsewhere (from 1986 to "Petroleumstilsynet"). The major off-shore accident in 1980 is therefore not in this data. We see from the graph below that the number of fatalities off-shore has been low in the period, and does not alter the trends above.

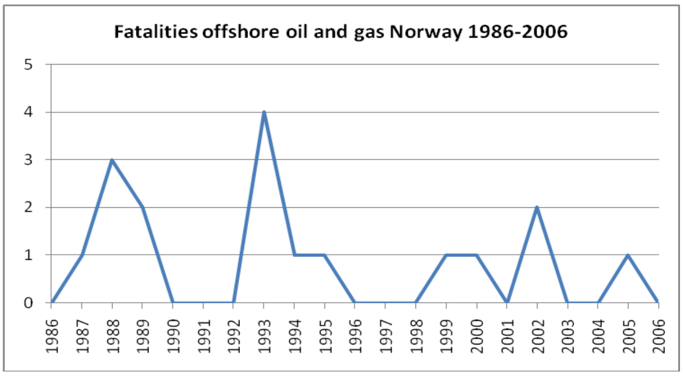

"Arbeidstilsynet" also reports each year the number of fatalities in two-way table, classified by 17 different occupations and 13 categories of causes. Similar statistics exists for the number of reported work-related injuries from 1990 to present. This shows a rise until 1998 and then a decline. However, injuries are generally underreported, and these trends may be due to changing reporting practices.

From the website of Statistics Norway, we may obtain various fatality and accident statistics, grouped by type of accident, region and gender. Here are the total numbers of pedestrian fatalities in Norway each year in the period 1996-2006: We see that there is a downward trend in the period.

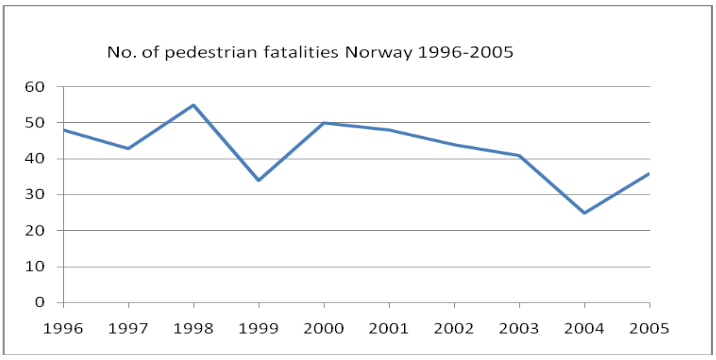

The statistics above is about yearly counts in a population, and do not tell that much about the hazards facing the individual. Human hazards may be judged by various measures, among them: Risk per hour, risk per year, risk per unit (different from time unit), risk per case and risk as shortened life-time.

<u>Risk per hour</u>: The hourly fatality risk for an activity is the average number of deaths per hour under influence of the activity, often reported per 100 mill hours.

Example: If the fatality risk in a profession is 1 per 100 mill hours, this may be interpreted as during 1 hour with 100 mill people at risk, one will die on average due to the activity. It may also be used to compute the yearly and lifetime occupational fatality risk of an individual in the profession, say by 1 250 hours and 50 000 hours respectively, assuming 40 years of employment. In the case of 1 fatality per 100 mill hours the yearly risk will be 1 in 80 000, while the lifetime risk will be 1 in 2000, which may be taken as among a workforce of 2000 people, one fatality is expected in their lifelong employment.

<u>Risk per unit</u>: The fatality risk per unit for an activity is the average number of fatalities per unit under influence of the activity, often reported per 100 mill units.

Example: In passenger transportation it makes sense to measure risk by fatalities per mill passenger kilometres instead of time, for instance when comparing different modes of transportation, e.g. flight and railway.

<u>Risk per case</u>: The fatality risk per case for an activity is the average number of fatalities per case under influence of the activity.

Example: What is the risk to go by plane to London?
Risks of this kind may alternatively be compared by the number of units required to give a specified fatal likelihood. For example: The number of units required to give a fatal risk of 1 to 1 million: 1 ½ cigarette for a smoker, 2 months for living with a smoker, 1 day for living in a New York pollution, 80 km for driving a car etc. Such statements are not so easy to interpret.

Expected shortened life: Is assessed with respect to exposure of a specific risk, say an occupation. It depends on at which age you start (left censoring) and how long you stay on (right censoring). The expectations in days may be given for different occupations, for being employed at different ages, for one year employment and employment to retirement.

The following compilation of hazards in various professions and other human activities in Norway is given by Elvik (TØI, 2005):

Activity	Fatalities per 100 mill hours
All work related (age 16-78, 2000-2003)	1.4
Agriculture and forestry	8.0
Fishery and hunting	10
Petroleum industries	3.3
Construction	3.0
Mining	15
Travelling on road (age 0-, 1998-2002)	18
Bus travel	3.3
Car travel	17
Motor biking (Large)	230
Motor biking (Small)	160
Moped	42
Cycling	23
Pedestrian	16
Railway travel	6.1
Air travel	34
Boat travel	7.6
Activities at home	1.7
Activities out of home	8.7

Many of these FAR values are down from those reported decades ago, due to increased awareness, improvements and regulations. In the 1970's when the offshore North Sea oil activities were in its infancy, the values in oil drilling was 20 and for the divers 100 (Samset: Teknisk Ukeblad, 1979).

Comparison of the risks for different professions, different modes of travel and different leisure activities is not simple, since we have to take into account the extent of hourly exposure. In land-based industries, this may be 7 hours a day 5 hours a week, but on a platform offshore you are exposed to some extent when you are off duty. On the other hand this may be compensated by longer periods ashore. When comparing different modes of travel, a higher risk for faster travel may be acceptable, since you travel far longer in an hour. Then the relevant figure may be fatalities per distance unit, say 100 mill km. It is even more difficult to find the proper basis for computing FAR numbers for risk sports, like mountain climbing and river canoeing. The FAR values may be high, but exposure hours are lower than most occupations. For some occupations, like professional boxing, the FAR values may be sliced by counting all the hours in the gym training.

Statistics saying "planes are safer than cars" are not likely to change people's behavior, due to habit, fear of flying, or love of driving. In fact, after the terrorist attack on September 11, 2001, the decision by millions to drive rather than fly, may have cost more lives than the 3000 lives lost at the World Trade Center.

Fatality statistics is mostly relevant at the national and supervisory level. For the individual enterprises, the fatalities occur so infrequent that it does not make much sense to track it. On the other hand, it will be important to track all accidents not leading to fatality and near accidents as well. The activity has to be precautionary so that accidents and fatalities do not occur. For this purpose it may be important to monitor key characteristics of systems and processes to see if changes occur that can lead into trouble, for instance by statistical process control charts.

It is now customary for enterprises to report performance indicators for HES in their annual reports. The indicators may be such as:

- Total recordable injury frequency per mill working hours (first aid excluded)
- Serious incident frequency per mill working hours
- Total sickness day as percentage of possible working days
- Spills and emissions
- Energy consumption
- Waste recovery factor

Here are two graphs from the 2006 Annual report of StatoilHydro, with totals including both StatoilHydro employees and contractors. Similar graphs are given separately for subsidiaries.

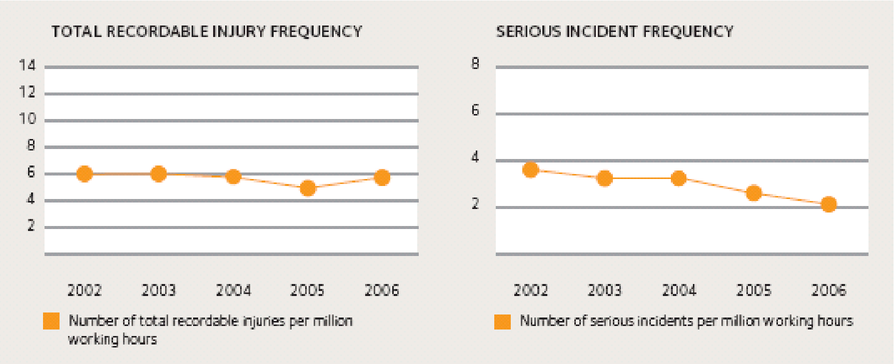

For judging injury statistics in general, the reporting habits may be a serious problem, as they may differ considerably among professions, locations and time periods.

Exercise
Examine the annual report of some major Norwegian companies, to see if and how they report HES related information for the year and make comparisons (trends etc.)

1.8 Accident investigations

Accident investigation may be viewed as four partly overlapping phases

1. Collection of evidence
2. Analysis of evidence
3. Drawing conclusion
4. Judgment of needs

There are essentially three levels of investigation according to frequency and severity of the incident (accidents and near accidents)

A. First level investigation
B. Problem solving group
C. Independent commission

At the workplace all reported incidents should investigated immediately at the first level by the supervisor and safety representative. For not very serious incidents, this may lead to some corrective action (in quality management terminology), and not lifted beyond this level. If the incident is serious, i.e. frequently recurring and/or high (but not extreme) potential loss (actual or possible), it should be handled by a problem solving group. Such incidences may have causes that are multiple, cross-functional and beyond the first level. Moreover, its prevention may require system changes. For these reasons, the member of the problem solving group should have diverse competence and sufficient authority. On rare occasions when the potential loss is very high, there is need for an investigation commission with independent status in relation to the organizations likely to be responsible for the incident. Typically this happens after a serious accident involving loss of lives, or potential loss of lives if repeated, although the probability is low. Whether this happened within a workplace with no outsiders affected or affected the public, like a plane crash, should not be decisive for having an independent commission. The context for these accidents is often modern technology, e.g. commercial aviation, petrochemical industry, nuclear energy plants. They may also have multiple causes, involving several organizational layers. Note that for big catastrophes caused by natural phenomena, like earthquakes, tsunamis and floods, it does not help much exploring the causes, but we still have the issue of how the consequences of the initiating event was handled and the preparedness for such catastrophes, e.g. the case of the tropical cyclone Katrina that hit New Orleans in 2005.

The accident investigations have to collect and analyze facts and statistics. The methods used may be diverse. In particular, the situations are very different when the data is scarce and data is abundant. Consider the following examples:

1. Road accidents
2. Safety at a workplace
3. Aircraft crash
4. Nuclear accident

The frequency of occurrence is here about inverse with the order. We have much data on road accidents, moderate on workplace safety, little on air crashes and very little on nuclear accidents.

This may call for different analytic methods. If we want to judge the effects of a specific measure, we also face problem of changes in other parameters in between scarce occurrences.

In relation to the discussion above it may be fruitful to think in terms of three safety control strategies (Rasmussen, Safety Science 1997):

I. Empirical Safety Control
II. Evolutionary Safety Control
III. Analytic Safety Control

Empirical safety control is to control the conditions and causes based on statistics from past accidents. This may work in situations like traffic accidents, infections at hospital and to some extent workplace safety. Evolutionary safety control is to control the accident process itself based on reaction to individual past accidents. This is the situation for aircraft crashes and train collisions. Analytic safety control is control of the accident process by predictive analysis of possible accidents. This is the only possibility for major nuclear and chemical hazards.

An overview of common methods for investigating major accidents, i.e. with emphasis on Evolutionary safety control, is given in "Methods for accident investigation" NTNU 2002.

In March 1980, the offshore platform *Alexander L Kielland* unit capsized on the Ekofisk field in the North Sea, with the loss of 123 people. How could this happen? Which actions and decisions, or their absence, led to such a tragedy? The nation was filled with disbelief, and everybody related to the business and their families was scared. The Norwegian petroleum industry would never be the same after this disaster. In July 1988, Britain's Piper Alpha platform exploded into a sea of flame, which killed 167 people. The investigations after Kielland and Piper uncovered a general failure to understand the factors, that may cause major accidents. These two disasters fundamentally changed the technical and organizational understanding of risk, the approach to petroleum activities, and changed the regulations, both in Norway and in the UK. This experience has also had impact to other parts of the world.

The understanding of risk in complex systems has also been enhanced by organizational studies and organizational theories. Key issues are: non-existent or degrading of safety culture, accumulation of latent conditions, organizational drift, lack of management commitment.

In Norway, the investigations after accidents in transport (aviation, railway, marine and road traffic) are handled by The Accident Investigation Board Norway (AIBN). ("Statens Havarikommisjon for Transport").

Exercises

1. How should we position an offshore oil spill in relation to the three categorizations (A, B, C), (1,2,3,4) and (I,II,III).
2. Give examples of publicly available investigation reports for categories I, II and III. Examine one!
3. Discuss the statement: "Safety is measured more by its absence than its presence" , (James Reason, 1997)

We may think of three kinds of threats to a system:

1. Regular threats – happened - recurrence always possible
2. Irregular threats – not happened - not likely – but imaginable
3. Unimaginable threats

Of interest here is the "Normal Accidents Theory" (NAT), due to the Yale sociologist Charles Perrow.[4] This theory essentially says that "No matter how hard you try, you will have accidents due to unexpected interactive complexity". Complex systems may typically exhibit behavior that was unanticipated at the design stage, and therefore conventional risk analysis methods have little predictive value. In Perrow's view, the two characteristics of a "normal" accident are:

- Unexpected and complex interactions between faults that are tolerable individually
- Tight coupling with little opportunity for mitigation or defense once a fault occurs.

His message is essentially that we cannot reliably predict the safety of a system on paper before it is built.

Here are some major incidents, with some links of varied nature and, not necessarily the most informative ones:

World	**Norway**
– Flixborough explosion, 1974 – Three Mile Island – Harrisburg, 1979 – Bhopal gas leak, 1984 – Chernobyl nuclear, 1986 – Challenger space shuttle, 1986 – Zeebrügge-ferry, 1987 – Piper Alpha explosion, 1988 – Exxon Valdez oil spill, 1989 – Estonia ferry, 1994 – Sea Empress oil tanker, 1996 – Longford gas explosion, Melbourne, 1998 – Tsunami, 2004 – Tropical cyclone Katrina, 2005 – Floods SE Asia, 2005 – Earthquakes, Turkey 1999, Iran 2004, Pakistan, 2005	– Bravo blow-out, 1977 – A. Kielland-platform accident, 1980 – Caledonien hotel fire, 1986 – Scandinavian Star ship fire, 1990 – Sleipner-platform, 1991 – Flood in Eastern Norway, 1995 – Norne helicopter crash, 1997 – Romeriksporten tunnel, 1998 – Sleipner-catamaran grounding, 1999 – Åsta-railroad accident, 2000 – Lillestrøm railroad incident, 2000 – Rocknes ship accident, 2004 – Hanekleiv tunnel, 2006 – Server ship accident, 2007 – Vest Tank explosion, 2007

An informative UK website is COMAH: Control of Major Accident Hazards.
Norwegian accident investigation reports are found at Accident Investigation Board Norway.

[4] Perrow (1984) Normal Accidents: Living with High Risk Technologies, Princeton University Press.

1.9 Risk control in municipalities

A municipality may face a number of hazards. Among them are:

Nature hazards	Human activity hazards
Extreme wind	Transport accident
Extreme rainfall (flood, intrusion,erosion)	Major fire
Landslide (rocks, clay etc.)	Handling of dangerous substances
Snow avalanche	Pollution
River flood (snow melting)	Critical infrastructures
Spring tide	Electromagnetic fields
Fire (wood, grass)	Terror and sabotage
Radon	Lack of emercency support

In the past the municipalities have mostly been reactive with respect to many hazards within their own responsibilities. As long as a project satisfied the formal requirements, there was often not much that could stop it. In recent years, the municipal responsibility to assess hazards has come to the forefront, among others after some recent major accidents. In addition to surveying existing hazards, the municipalities should incorporate safety and precautionary elements in the ordinary planning process.

Example: Building permits given for residences where they should not have been:
– Floods (Eastern Norway, 1995), Landslides (Hatlestad, 2005), Rockslide (Ålesund, 2007)

According to The Civil Protection Act (”Sivilbeskyttelsesloven”, 2010, §14) the municipalities have the responsibility for preparedness by doing risk- and vulnerability assessments. Efforts have to be made to reveal potential adverse events to the public within its borders, assess the likelihood of such events and their effects. The output of this work should be put together in a comprehensive risk- and vulnerability report, which should be the basis for the work for security and preparedness to the benefit of the public.

Enterprises and developers in Norwegian municipalities have to comply with a number of laws and regulations with respect to hazards and safety. Central to this is The Planning and Building Act (“Plan- og bygningsloven”, 2008), requiring that developers include a risk- and vulnerability assessment in their project plans in order to get approval from the authorities. For the sake of societal safety, hazard spots and vulnerable spots should be marked in area plans as zones requiring special consideration. The different regulations may be supervised by the municipality itself, the county or some specific national body. Of particular interest is the role of the county governors (“Fylkesmannen”), who is the State representative office in each county. In recent years county governors have frequently pointed fingers to the communalities for their lack of proper risk- and vulnerability assessments, and have in some cases stopped projects.

Regulations warranted in law typically belong to a specific ministry, which may have a directorate or supervising authority for a given area of concern, responsible for more detailed regulations and the supervision of these regulations. Some directorates of relevance to municipalities are the following:

- Directorate for Civil Protection and Emergency Planning under Ministry of Justice
  - In Norwegian: Direktoratet for samfunnssikkerhet og beredskap (DSB) Net-address: www.dsb.no
- Norwegian Water Resources and Energy Directorate under Ministry of Petroleum and Energy
  - In Norwegian: Norges vassdrags- og energidirektorat (NVE), Net-address: www.nve.no
- The Norwegian Directorate of Health under the Ministry of Health and Care Services
  - In Norwegian: Helsedirektorat Net address: www.Hdir.no
- Directorate of Public Roads under Ministry of Transport and Communication
  - In Norwegian Vegdirektoratet Net address: www.vegvesen.no
- Norwegian Environment Agency under Ministry of Environment
  - In Norwegian: Miljødirektoratet Net-address: www.miljodirektoratet.no

These are executive agencies on the national level with competent authority. In addition there are supervising authorities. Among them are

- Norwegian Labour Inspection Authority under Ministry of Labour
  - In Norwegian: Arbeidstilsynet Net-address: www.arbeidstilsynet.no
- Norwegian Board of Health Supervision under Ministry of Health and Care Services
  - In Norwegian: Statens Helsetilsyn (Htil) Net-address: www.helsetilsynet.no
- Norwegian Radiation Protection Authority
  - In Norwegian; Statens strålevern (NRPA) Net-address: www.nrpa.no
- Petroleum Safety Authority (PSA) under Ministry of Labour
  - In Norwegian: Petroleumstilsynet (Ptil) Net-address: www.ptil.no

Take the case of building projects. This may range from individuals asking for building permits with limited consequences to investors wanting to locate a new factory with potential hazards far beyond the enterprise itself. In the first case the municipality will follow standard procedures and regulations: Control that the plans are acceptable and the neighbours are given notice prior to giving building permit. It should also see that the building is according to common building standards with respect to dimensioning and use of materials, and that the construction takes place without hazards to the workers and the neighbourhood. In most cases this is handled by the municipality, but the county (and in some cases a ministry) may interfere or be the body for appeals. In the latter case it is not just an administrative issue: It is beyond the daily routines, the stakeholders are many, in many cases the public at large, so the issue is political. It also requires competence not readily at hand in the municipality. Ideally a risk analysis should be undertaken and made public at an early stage, before the project comes to the point of no return.

A simple form of risk analysis, named ROS-analysis, is recommended by the Directorate for Civil Protection and Emergency Planning (DSB) and the counties, although not explicitly warranted by law. The objectives of ROS-analysis are (Norwegian: ROS=”Risiko og Sårbarhet”=”Risk and Vulnerability”):

- Provide an overview of community hazards and vulnerability
- Evaluate consequences and likelihoods
- Identify and rank risk reducing measures

A good ROS-analysis should comply with the standard NS-5814 mentioned in section 1.6).

A simplified version of ROS named DagRos ("Daily ROS") is published by the Norwegian Directorate of Health and Social Welfare, and is suitable for simple everyday situations and limited context, like a kindergarten field trip. (See English brochure Mini Risk Analysis).

The Directorate for Civil Protection and Emergency Planning (DSB) provides guidelines on how to do a ROS-analysis in a municipality.

A summary sheet like the one below is suggested, with some guiding on how to fill the boxes. The grading of likelihood and consequences is mainly done by marking the statement that fits best, without stating exact numerical values. .

<table>
<tr><th colspan="3">Registration of risk and vulnerability</th></tr>
<tr><td>Description of adverse event</td><td></td><td>Supporting documents</td></tr>
<tr><td>Cause of event</td><td></td><td></td></tr>
<tr><td>Proactive Controls</td><td></td><td></td></tr>
<tr><td>Likelihood</td><td><table><tr><td>Very unlikelely</td><td>Somewhat likely</td><td>Likely</td><td>Very likely</td><td>Not assessed</td></tr><tr><td></td><td></td><td></td><td></td><td></td></tr></table></td><td></td></tr>
<tr><td>Reactive Controls</td><td></td><td></td></tr>
<tr><td>Consequence Description</td><td></td><td></td></tr>
<tr><td>Consequence grading for...</td><td><table><tr><td></td><td></td><td>Not dangerous</td><td>Some danger</td><td>Critical</td><td>Dangerous</td><td>Cata-strophic</td></tr><tr><td></td><td>Man</td><td></td><td></td><td></td><td></td><td></td></tr><tr><td></td><td>Environment</td><td></td><td></td><td></td><td></td><td></td></tr><tr><td></td><td>Econ. value</td><td></td><td></td><td></td><td></td><td></td></tr><tr><td></td><td>Op./Prod.</td><td></td><td></td><td></td><td></td><td></td></tr><tr><td></td><td>Other</td><td></td><td></td><td></td><td></td><td></td></tr></table></td><td></td></tr>
<tr><td>Risk</td><td></td><td></td></tr>
<tr><td>Suggested Actions</td><td></td><td></td></tr>
<tr><td>Remarks</td><td></td><td></td></tr>
<tr><td>Prepared by</td><td colspan="2"></td></tr>
<tr><td>Date</td><td colspan="2"></td></tr>
</table>

Here follows an outline of the planned work by a risk- and vulnerability working group in Riskstad, a fictitious large municipality with a varied risk picture.

Example Municipality

Municipality:	Riskstad
Project:	Risk- and vulnerability analysis
Steering group:	Head of administration. + 4 additional
Working groups:	Electricity supply, water supply, transport of dangerous material + group for hazards for humans, environment and society in general
Working group 3:	Transport of dangerous material
Group expertise:	Technical, environment, medicine
Information sources:	Experts, locals, locations, activities, (near) accident statistics, supervisory reports, existing routines
Adverse event 1:	Release of dangerous chemicals
Specific Information:	Road, traffic and railroad authorities
Possible causes	Human, technical, organizational, outside hits (man, nature)
Preventive measures:	Proper localization of facilities, transport routes, active maintenance, detectors and alarms, trained employees, protection, containment
Likelihood:	Unlikely, Not very likely, Likely, Very likely (meaning?)
Consequences for:	Man, environment, economic values
Consequence scale:	Safe, Some danger, Critical, Dangerous, Catastrophic (meaning?)
Summary:	Tabular
Initiative:	Preventive/containment, human/technical, responsibilities

Agenda
1. Map adverse events
2. Causes and likelihood
3. Consequences
4. Summary
5. Initiative

From the description of ROS-analysis it is clear that this is insufficient as basis for decisions regarding situations that may have major hazard potential. It may be useful to get an overview of existing hazards in the municipality, but not sufficient to evaluate new projects with hazard potential. Developers typically have to provide a risk analysis to the municipality in order to get acceptance for such a project. At the early stage, long before final decision of acceptance is made, a simple analysis may be sufficient. A ROS-analysis may then be used as a preliminary tool to see what kind of extensive risk analysis that may be required. But be aware, developers may benefit to do as little as possible in this respect and also play down the risk of the project. By accepting just a simple analysis, the municipality and the politicians may risk that the project goes beyond the point of no return without satisfactory evaluations of the risks involved.

In every land-based activity handling poisonous or dangerous chemicals, accidents may happen which may lead to dramatic consequences for material values, humans and the environment. Accidents may happen in production, storage and transport at the location. The risk of accidents by release of dangerous chemicals, and the risk of subsequent fire and explosion should be addressed, and be reduced by establishing appropriate precautions. This applies of course also to projects under way, and before decisions to accept such activities at the location.

Example: The LNG- project in Risavika 2005-2010

- A ROS-analysis was made at the preliminary stage.
- Complaint to the county governor, but ROS ok'ed
- Extensive risk analysis overlooked major scenarios?

– The role of the consultants: Provided just what was asked for?
– The role of local civil servants and politicians: Conflicts of interest?
– The role of DSB questioned: Competent or competence not used in the public interest?
– Can the follow-up process remedy the initial weakness?

What are acceptable individual risks in the setting of projects affecting third parties? Numbers frequently referred in practice are the following

Risk category	Individual risk (fatal probability per year)	
	First and second person	Third person
Unacceptable	$> 10^{-3}$	$> 10^{-4}$
ALARP	10^{-6} to 10^{-3}	10^{-6} to 10^{-4}
Acceptable	$< 10^{-6}$	$< 10^{-6}$

In the case of major hazard potential involving chemicals and the danger of fire and explosion, the so-called Seveso directive applies (named after the incident in Seveso, Italy in 1976). This directive was issued by the European Commission, dating back to 1982, with revision Seveso II from 1996 and later amendments, aiming at

> "…. the prevention of major accidents which involve dangerous substances, and the limitation of their consequences for man and the environment, with a view to ensuring high levels of protection throughout the Community in a consistent and effective manner".

The directive states an obligation for companies to provide information on the risks that their operations pose to the local population. The EU member states have made efforts to adopt regulations that comply with this EU-directive, with names such as "The Hazards of Major Accidents Decree" (In Norway: "Storulykkeforskriften", administered by DSB). According to this the responsibility to report is limited to specific dangerous substances and the quantities handled. It is also limited to specific large enterprises.

Example The Vest Tank accident
– Storage of oil and chemical waste at Sløvåg blew up May 2007
– Vest Tank not under the reporting decree
– Competence: Did they know what they handled?
– Crisis management: The municipality, SFT, DSB
– What can be learned?

The municipalities should preferably have or develop plans for crisis situations, where ordinary routines and resources are insufficient to handle the situation. Of course the municipality will get help from outside, the county or from national level. However, in some cases, the time factor, confusion and unclear responsibilities, both on local and national level may delay prompt action.

Exercises
1. Find a more detailed description of ROS-analysis for communalities (**http://www.dsb.no**/)
2. Find examples of various risks in a communality for different combinations of
   (i) perceived/not perceived (ii) happened/not happened (iii) ignored/addressed
3. Perform a Mini Risk Analysis for a kindergarten field trip (See brochure)

1.10 Societal security

The term societal risk has materialized in recent years, meaning risks to the public at large or the society itself. However, there may be disagreement on how wide this notion should be applied. Some possibilities are:

1. Infrastructure breakdown, inherent to systems.
2. Attacks on systems from outside
3. Pandemics
4. Endangered living condition due to environmental changes

Some prefer to limit the notion to risks specific to living in a modern complex society, i.e. to infrastructure breakdowns and attack from outside (hackers, terrorists etc.), leaving out pandemics and environmental change, but include bio-terrorism. Anyway, it seems natural to leave out the risks related to military aggression, risks from inside the society due to breakdown of citizen morale or cooperation. It also excludes the common local or regional environmental risk, due to drought, floods and earthquakes. In some cases it may be difficult to draw a separation line. For societal risks, we expect that the government is a provider of societal security, another term come to use (not to be confused with social security).

The societal risks listed above are very different. What they have in common is that they are beyond local authorities to handle, with respect to preparedness, required immediate actions and handling of consequences. In general we should of course strive for an anticipative society, but also be better at building resilience into our society, i.e. capability to absorb and recover from adverse extreme events. Instead of putting more efforts into gaining more knowledge on uncertainties, it may in some cases pay off to study how to live with uncertainties and be adaptive to surprises at the same time.

Societal risk and security is a field too broad to discuss in any detail here, and we will below limit ourselves to some selected issues related to the first issue listed above and how this may be handled, and then only briefly open up for discussions on the other issue.

Risk in socio-technical infrastructure systems

A modern society is heavily dependent on large socio-technical service systems. They shall provide prosperity, security and general wellbeing for the citizens. Among the vital infrastructure systems are the utilities, like electricity and telecommunications, and the transport systems. Breakdown and disruption of such systems may have severe consequences, and even small disturbances may be very annoying to the public. The causes of such disturbances may be diverse:

(i) Natural disasters, like earthquakes, floods, hurricanes or other adverse weather conditions.
(ii) Technical failure, most often with a human factor as roots cause, in planning, operation or maintenance.
(iii) Criminal acts, like sabotage or terrorism.

The following scheme is used in an official report (NOU 2006:6) on the protection of critical infrastructures and critical societal functions in Norway:

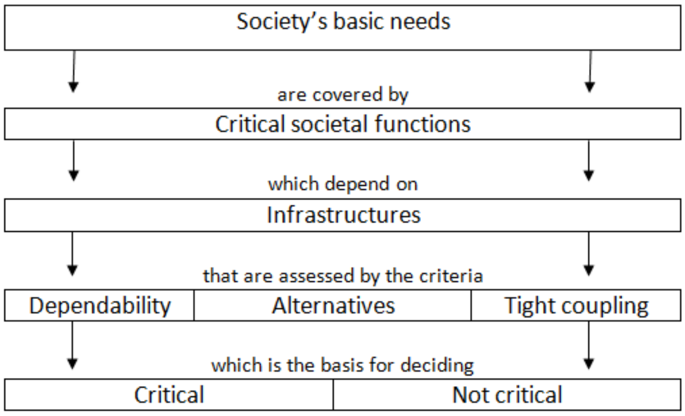

Here criticality is implied by any of

- high degree of dependability
- few or no alternatives
- high degree of coupling, e.g. in a network

The following critical societal functions and critical infrastructures are identified and dealt with in some detail in the report:

Critical societal functions
Banking and finance
Food supply
Health services, social services, social security benefits
Emergency and rescue services
Waste treatment
Environmental surveillance
The police

Critical infrastructures
Electrical power
Electronic communications
Water supply and sewage
Transport
Oil and gas
Satellite based infrastructure

In particular, the role of public ownership is discussed, when relevant.
The parliament, government, judiciary and the defense add to the list of critical societal functions, but they are not dealt with in any detail.

Socio-technical systems may be studied along three main dimensions (Linstone, 1984):

- The technical dimension
- The organizational dimension
- The personal dimension

In addition to this, there may be dimensions like geographical, economical and institutional.

We have essentially three ways of assessing the likelihood of an adverse event: (i) by observed frequency of previous occurrences, (ii) by mathematical modeling using experience of component failures and (iii) by expert judgment. The opportunities for assessing the consequences may be divided along similar lines.

Typical for most socio-technical systems of today is the lack or incompleteness of relevant breakdown accident data. This is partly due to lack of real (worst scenario) incidents, and that the system change over time. Some examples: For explosions at a chemical plant we may have models for the diffusion of gases, but perhaps no knowledge on long term effects and the environment. For electric power outages, we may have data on frequency, unserved energy, power loss and restoration time, but scarce data on the total effects for enterprises, services and the public at large.

For risk assessment of socio-technical systems, the following may be quantities of interest:

- Occurrence of disturbance or breakdown (measured by frequency or described by the occurrence process itself)
- The severity of disturbances
- The time to restoration

For this, we may need statistical modeling tools in combination with knowledge of the typical features of such quantities. Some distribution theory and process theory may be helpful.

A different challenge is the assessment of the risk potential related to planned attacks. Following traditional risk analysis we have to assess the likelihood of various types of attack and the likelihood of various adverse consequences, given that the specific type of attack has taken place. The hopes are that this provides an overall risk picture, and guide the kind of measures required to protect from attacks and/or their consequences. However, there is a weakness in this way of thinking. The potential attacker may learn about the specific measures taken, and direct their efforts according to this knowledge. The situation then has some resemblance to a game between adversaries, and perhaps game theory may be helpful for this kind of risk analysis.

Discussions:

1. Discuss societal risk in the following areas:
   (a) Utilities (b) Production (c) Transportation
2. Are we good at thinking the unthinkable?

Discussion: Electric power supply
The most important infrastructure system of modern society is probably that of electric power supply. What happens if the energy supply fails:

— On Thursday December 21 at 0900 am.
— Children still at school and kindergarten
— It is rather cold (-10°C)
— The Christmas vacation travelling started

Let us briefly focus on some issues related the last three points listed in the beginning of this chapter, which may also be a platform for interesting discussions:

IT- failures and attacks

Information technologies have experienced tremendous improvements in speed, capacity and opportunity over the last decades. The products and services have also become more user-friendly with less downtime. However, many systems we all depend upon have become increasingly complex, and may fail due to its inherent complexity or due to attacks from outside. Both may be difficult to handle without serious harm to individuals, companies and the society. For this reason, diagnostic systems and "firewalls" are made, but it is hard to prepare for every eventuality, and some protections may even add to the complexity of the system. Several developments in the last decade have to be of concern, among them the appearance of the internet and software integrated with the operating system. The combination of these two, where software is automatically upgraded over the net seems extremely dangerous, from both the viewpoint of complexity, and the hacker threat. This was something not taken very seriously in the beginning, where the focus was to make firewalls to stop hackers to the system.

Discussion: How could a IT-failure or attack affect you personally?

Security technology

Various technologies have come to use In order to improve societal security. Among these are: Sensors, communication technology, data storage, analysis and decision support systems. Common for many of these is that they may intrude privacy and may be abused.

Discussion:

1. Is privacy a threat to security? How to balance? Who should decide?
2. Is there need for regimes where: "You are guilty until proven otherwise"?

Bio-surveillance

Bio-surveillance is the monitoring of selected information sources for early warning of emerging epidemics, whether it be naturally occurring or due to of bioterrorism. Indicative information may be diverse: Increased purchases of nonprescription medication, increase in reported symptoms during ambulatory care, and preferably the reporting of diagnostic results confirming the presence of a pathogen. Such surveillance may routinely be performed by public health authorities. A workable system requires: Relevance, frequency, timeliness, sufficiency and accuracy. With most surveillance systems based on statistics, we have the problem of false positive. To face a variety of possible threats, some even unknown, is a difficult task!

Discussions:

1. How could a pandemic affect society?
2. What are the hard decisions for the authorities with respect to preparedness?
3. What are the challenges with respect to informing the public?

Environmental change

The major thinkable changes relevant to Norwegians are:

1. Collapse of “Golf-strømmen” (The Atlantic Meridional Overturning Circulation) .
2. Melting of the North Pole ice cap.
3. Global warming with sea level rise.

Discuss the possible consequences of each of the listed environmental changes

The following may be some broad categories of causes of societal risk in a modern society:

1. Complexity: Mismatch between the system and the available mechanisms to control it.
2. “Flat world” instability: Due to instantaneous global access to the same information, leading to synchronous acts (herd and bubbles), instead of asynchrony acts with balancing effects.
3. Paradigm shifts: Discontinuity of socio-economic trends globally or regionally, e.g. with respect to technology.
4. Maturity of global crises: A long term adverse global trend reaching a tipping point, possibly of no return, caused by lack of global governance.
5. Built-in myopic features: Major driving forces are individualistic and short term oriented, and no one takes a long term view or has the power to counterbalance the short term view.

Discussions:

1. To what extent each of these applies to the four topics listed in beginning of the chapter.
2. The number of natural disasters and their losses have apparently increased in recent decades. Give a possible explanation.
3. Discuss how poverty and affluence might influence the vulnerability to hazards of Nature.
   - You are told that deaths from natural disasters have decreased in developed countries and increased in developing countries. What could explain this? Politics? Economics? Education? Culture?
   - You are told that costs of natural disasters have increased in developed countries. What could explain this?
4. Discuss whether the following may increase societal risks and threats:
   (a) Globalization, (b) Privatization, (c) Deregulation, (d) Urbanization, (e) Mass travel

In a recent book [5] Perrow (2007) argues that the high risk in our modern complex society is linked to the following three main factors:

1. Concentration of humans
2. Concentration of energy
3. Concentration of corporate power

In this book he goes beyond the risks “designed into the system”, and also discusses the risks associated with unfriendly attacks on the system from outside. His main message is that it is not sufficient to protect against or reduce damage. We should *Reduce the targets!*

[5] Perrow (2007) The Next Catastrophe: Reducing Our Vulnerabilities to Natural, Industrial, and Terrorist Disasters. Princeton University Press

Discussions:

1. Discuss the relevance of Perrow's three risk factors for societal risk.
2. Discuss the objective "reduce the targets" in relation to various disasters: hurricanes, floods, explosions, fires, utilities breakdown, It-collapse (computers, databases, telecom etc.)

Authorities

For accidents within a Norwegian municipality, the municipality authority itself is responsible for the first line response, except for certain accidents where another pre-assigned authority is responsible, e.g. salvage operations when a ship is grounded. The local authority may have to ask for assistance from the county governor's office and national authorities, but still being responsible for the operations. In some recent emergency situations, some municipalities have felt that central authorities have not provided prompt and adequate help, even when asked for, perhaps due to unclear responsibilities on their side. This may of course happen in situations when something unimaginable crossing several areas of competence has happened, but should not be so.

The main central authority in this area is the Directorate for Civil Protection and Emergency Planning ("Direktoratet for Samfunnssikkerhet og Beredskap", DSB). The objectives are:

- Provide one joint line of authority from central to local level within the areas of fire, rescue and general preparedness
- Maintain an overview of risk and vulnerability for the Norwegian society in general, and provide professional expertise covering prevention and preparedness for incidents at central, regional and local levels.

The challenges for DSB range over a wide spectrum, from ensure that vital public functions are not paralyzed to provide information to the public regarding their own safekeeping, including safety aspects of marketed products and of consumer services.

Vision: "A safe and robust society where everyone shares the responsibility to safeguard life, health, the environment, vital public functions and material assets".

Details on DSB activity may be found from the web-pages http://www.dsb.no/.

Besides this, there is also The Norwegian National Security Authority (NSM), which is a cross-sector professional and supervisory authority within the protective security services in Norway. The purpose of protective security is "to counter threats to the independence and security of the realm and other vital national security interests, primarily espionage, sabotage or acts of terrorism. Protective security measures shall not be more intrusive than strictly necessary, and shall serve to promote a robust and safe society". It is under the Department of Defense, but report to the Department of Justice in civilian matters. (www.nsm.no).

Discussions:

1. "Norway is considerably more vulnerable than before"
2. Is the local responsibility model the best?

To broaden the scope to the international arena, it may be of interest to examine material from the International Risk Governance Council. IRGC is an independent non-profit organization founded in 2003 and based in Geneva Switzerland, with mission to improve the understanding and management of global risks with impact on human health and safety, the environment, the economy and society, at large. From IRGC material: "We work to achieve this mission by reflecting different views and practices, and providing independent, authoritative information, by improving the understanding and assessment of risk and the ambiguities involved, by exploring the future of global risk governance and by designing innovative governance strategies." Among others, IRGC has developed "an integrated analytic framework for risk governance which integrates scientific, economic, social and cultural aspects and includes the effective engagement of stakeholders". The main focus, is on low-probability high-consequence outcomes of wide-range concern, and goes beyond most common risk management frameworks, dealing with how risk-related decision-making processes typically unfold, with the need for coordination and reconciliation between stakeholders. For more, see http://www.irgc.org/.

The United Nations has established International Strategy for Disaster Reduction (ISDR) with office in Geneva. Their mission is:

"The ISDR aims at building disaster resilient communities by promoting increased awareness of the importance of disaster reduction as an integral component of sustainable development, with the goal of reducing human, social, economic and environmental losses due to natural hazards and related technological and environmental disasters"

ISDR promotes four objectives as tools towards reaching disaster reduction for all[6]:

- Increase public awareness to understand risk, vulnerability and disaster reduction globally
- Obtain commitment from public authorities to implement disaster reduction policies and actions
- Stimulate interdisciplinary and intersectoral partnerships, including the expansion of risk reduction networks
- Improve scientific knowledge about disaster reduction

For more on ISDR see http://www.unisdr.org/.
Societal and global risks may differ with respect to the level of knowledge. Different levels of knowledge require different strategies for risk treatment, that is, whom to involve and what to do. The main strategies are often labelled risk-based, cautionary, precautionary and discoursive (i.e. inform, build confidence, involve the affected). Many situations require a combination of one of the first three with the latter one. More details on risk management strategies may be found in section 2.5.

[6] For details see UN/ISDR: Living with Risk: A global review of disaster reduction initiatives 2004 version

1.11 Risk and the public – perception and communication

When a person is given a list of a variety of human fatality risks, and is asked to rank them according to risk level from high to low, it is well known that the ranking may be wildly different from the ranking from actually observed data, even if the outcome (death) is the same in all prospects. Even when available "objective data" tell that some perceived high risks have probabilities of a much lower magnitude, this may not change the opinion. Reasons for some risks to appear more threatening than others may be:

- they are felt to be completely out of our control,
- they happen instantly, sometimes spectacularly.

A notable example may be the comparison of travelling by own car and travelling by airplane. In the latter case, we are completely dependent on the traffic monitoring system and the flight crew, but in the former case we feel we have control, even if we are dependent on the state of the car, road conditions and the traffic (reckless drivers etc.). Other threatening situations felt to be out of our control are discovery of E.coli bacteria infected food at a meat packaging plant, and reported incidence of Creutzfeld-Jacob disease, even if this may be an isolated case abroad. Some real threatening, but not very likely scenarios, are being affected by pandemic flu or a terrorist attack.

On the other hand, we have the fatal long-term risks, quite often related to lifestyle and habits. In some cases, the effect is not settled and different views exist among experts. The media thrive on both the spectacular and the long-term lifestyle risks. The public may easily be bewildered in both cases. Let us discuss some aspects of this.

Lifetime risk and the media

The media frequently make headlines implicating that some food, beverage or habit is dangerous to your health. For example

"xxx causes cancer"

In some cases this is accompanied by specific statements of recommendations like

> "Scientists recommend that you stay away from xxx" or
> "Scientists recommend that your intake of xxx is less than yyy per week"

The media will of course be happy if this provokes some reaction from someone who questions the conclusion or recommendation. In some cases the conclusion may of course be flawed due to scarce data or misinterpretation of the data, but suppose that the science regarding the causality is essentially valid. Next, the recommendation may be ill-conceived, due to scientists with too narrow scope and outside their responsibility. However, quite often the scientists have not made any specific recommendations for the individual at all. The headline is created by an eager journalist (or the desk) who wants attention. This is done by

- reading the scientific report very selectively,
- misinterpret or exaggerate the message,

- confuse the individual perspective vs. the population perspective

We will look at the latter.

Research often reports relative risks, saying that people in one group with a specific habit have a significant larger risk of arriving in some unwanted state (say cancer) with shortened lifetime than another group. A question may be whether significant here means just statistically significant or that the difference is of practical significance. Science usually reports statistical significance[7], and leave to others (politicians or regulators) to judge whether it is of practical significance and what to do about it. You may have statistical significance with no practical significance, since with abundance of data you can discover even minor differences. You may have both statistical significance and practical significance, but then the question may be: For whom? You should not allow to be fooled by a large relative risk, without asking yourself what is the absolute risk. Maybe the absolute risk is so low that it feels of no relevance to you, and you do not want to give up a habit you like.

Here the issue of individual perspective versus the population perspective comes into effect. Suppose the majority of the population is low risk. For the community it may pay off if it gets sufficiently many to adjust their habits, even if the effect is minor for each individual; this instead of trying to identify the few high-risk individuals. Therefore, it happens quite often that the public complains about "The nanny state", when singled out and seemingly unnecessary advice is given. This is sometimes referred to as

The prevention paradox: "A preventive measure that brings large benefits to the community and offers little to each particular individual" (Rose, 1992).

Nobody is "right" or "wrong" in this debate. It is perfectly rational for the government to give general advice based on the population perspective, and it is perfectly rational for the individual to ignore it based on the individual perspective.

Quite often issues of causation are debated in the media, and the debate is often clouded by ideologies and emotions. Take the case of violence on TV. Some people argue that violence on TV must be censored because it causes people to like violence, while other people argue that there is violence on TV because people like violence, and see no need of censoring, since violence exists in the world in the first place. In this example, it is not obvious what the cause really is and the issue maybe clouded by high emotions.

Statistical information has been available to the public since the early 1900's and some places even before. However, information about the state (and therefore named statistics) was many places regarded as state secrets, and availability to the public came with the rise of the democracies. Although statistics are public, the public is not very numerate.

For more on perception of risk and causation and risk literacy, see section in Part 2.

---

[7] The practice of only reporting studies with significant differences may lead to biases in the accumulation of knowledge.

Crisis management and communication

The disparity between perceived risk and "objective risk" provides a challenge for public servants and business management, when something or this kind happens. This challenge may be named *crisis management and communication*. The disparity also gives an opportunity for interest groups and the media. This may be utilized both prior to and after any hazard, with emphasis on the spectacular incidences and long-term scenarios, out of our control, but not very likely. This use may of course be justified in order to awaken the public and/or find and challenge someone responsible, but may also be misused to gain attention in order to suit their own interests. Finally, the disparity gives an opportunity for consultants. Such consultants may be on both sides, for instance: (i) To help uncover and establish awareness of an environmental danger and (ii) to help a business to handle the communication to the public after an environmental incidence has occurred, like a major oil spill. It is of interest to see how risk consultants of this kind may view their field.

The influential risk consultant Peter Sandman[8] reduces his field to the formula

$$\text{Risk} = \text{Hazard} + \text{Outrage}$$

Here "hazard" means a mode of death and its attached death rate, and "outrage" means everything else that the public attaches to the hazard. His emphasis is on outrage, i.e. the public reactions to the hazard.

Some dimensions of outrage are

- voluntary/non-voluntary
- familiar/unfamiliar
- detectable/undetectable
- immediate/delayed
- not memorable/memorable
- diffuse in time and space/focused in time and space
- fair/unfair
- controllable/uncontrollable
- controlled by the individual/controlled by others
- imposed by institutions that are trustworthy/untrustworthy

A risk may be termed high, if it is high in some of these dimensions, even if the hazard is low.

The field may be divided in three as follows:

A. Crisis communication (High Hazard, High Outrage)
B. Precaution Advocacy (High Hazard, Low Outrage)

[8] See http://www.psandman.com

- among others Safety and Health communication

C. Outrage management (Low hazard, High Outrage)

To reduce the public concern about a small hazard, the risk manager has to diminish the outage. Sandman points out six strategies for managing outrage:

1. Stake out the middle not the extreme
2. Acknowledge prior misbehavior
3. Acknowledge current problems
4. Give others credit for achievements
5. Share control or be accountable
6. Bring unacknowledged concerns to the surface

We probably have to accept that experts and public servants have far from perfect knowledge on the most likely outcome of different choices of action, as well as the public preferences for the outcomes. However, exposing uncertainty is often taken as insufficient work or insufficient competence, in particular in the media. Definite answers are wanted! Dialogue in some form between the stakeholders is required. Risk needs to be understood by all stakeholders, and not the least, the media.

Exercise

News media often report the number of accidents of a specific type last year and compare it with the preceding year. Example: "*The 18 accidents last year is up from 12 the preceding year. This 50% increase is a serious setback*". State some critical comments to this reporting.
Suppose you figure out that the yearly average over the ten preceding years was 16 accidents. Add to your comment, based on your judgment on what is unexpected and deserves special attention. What kind of additional information could be useful?

2 Approaches and tools for risk management

2.1 The risk management process

The main steps were as outlined in section 1.2 (Clause 5 of the ISO 31000: 2009)

1. Establish context
2. Risk assessment (Risk identification, Risk analysis, Risk evaluation)
3. Risk treatment

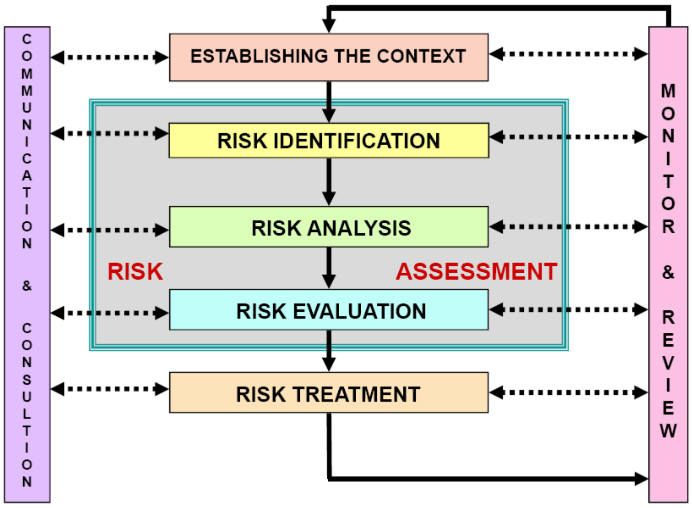

<u>Establish context</u>**:** Questions to be asked are:

– What is the activity to be studied?
– What is the aim of the study? Who are stakeholders?
– What attributes (e.g. life, health, economic value, environment) shall we study?
– What are the criteria for judging the risk?
– What are the opportunities and limitations?
– What kind of risk analysis suits the context?
– How will the results from the study be used by the stakeholders?
e.g. in daily operation, for strategic decisions or for political activities

This may be regarded as the planning phase for a risk analysis. It is wise to spend substantial time on this phase and address the questions in a systematic manner. Otherwise some confusion will typically remain on what we are going to accomplish, and time will be wasted later on to sort this out.

Risk identification and risk analysis: The purpose is to create a proper understanding of the risk elements in the activity, process or system in question, its aims and role for the stakeholders, including strategies and ramifications. This understanding should be presented in a form understandable for the stakeholders. The three main analytical tasks of a risk analysis will be

- Identify potential (significant) risk modes

- Identify components, systems, activities and processes that affect the identified risks.
- Identify and compare different alternatives with respect to reducing or removing the identified risks, where risks and costs are jointly judged.

A risk analysis may be performed when planning products and systems, activities and processes. It may also be required during operations for improvement or handling unforeseen problems, and even when closing down an activity. Today many businesses and face some regulations related to risk which are supervised by some authority. A risk analysis may then be required to document safe operations according to the regulation. It is convenient to imagine three levels of risk analysis:

(a) Simplified (qualitative)
  - Informal by just brain storming and group discussion
  - No formal risk analysis method
  - Risk measured on crude scale (Low, Moderate, High)

(b) Standard (qualitative or quantitative)
  - Formal systematic methods used

(c) "Advanced" (mostly quantitative)
  - Model based risk computations
  - Event trees, fault trees etc.
  - Probabilistic and statistical modelling

Parts of creating the understanding of the activity in question is common to all three levels of risk analysis, and by some authors regarded as part of creating the context Ingredients to look for are the 5 M's: Man, Machine, Material, Method, Milieu and their relationship, typically viewed and described within a process framework. The creation of understanding should preferably be performed as a creative process in a diverse group, tapping the knowledge of everyone in the group. Visual techniques like flow diagrams may be helpful for creating a common understanding and be the basis for later analysis. It is dangerous to neglect the description stage, since lack of understanding of this is often prevalent among those working in the system, which may itself be a source of unwanted risks. A risk analysis typically involves collecting, digesting and presenting data. However, the presentation of historic data does not constitute a risk analysis per se. A risk analysis should always have a predictive view. Presentation of historic data without reflections on whether they are relevant for the future is not sufficient!

A typical context for risk analysis is where we have an initiating event or incidence (B) followed by an outcome or consequence (C). Prior to the initiating event, we may have some possible causes (A), so that we imagine

$$A \rightarrow B \rightarrow C$$

,: The objective of this step is to determine whether some of the risks uncovered by the risk analysis need to be acted upon, and set priorities. We then have to compare the risk levels with the risk criteria determined at the first step of establishing context. If the uncovered risks do not meet these criteria, they have to be treated. By evaluating the disparities, we can set priorities. However, cost-benefit analysis or other relevant criteria may also lead to the conclusion that no risk controls should change. It is important to pay attention to rare and extreme risks which are not easily or possible to cast in economic framework.

A so-called Bow-Tie diagram is frequently in use in practice for keeping the context for risk analysis in the minds of the participants. It looks like this:

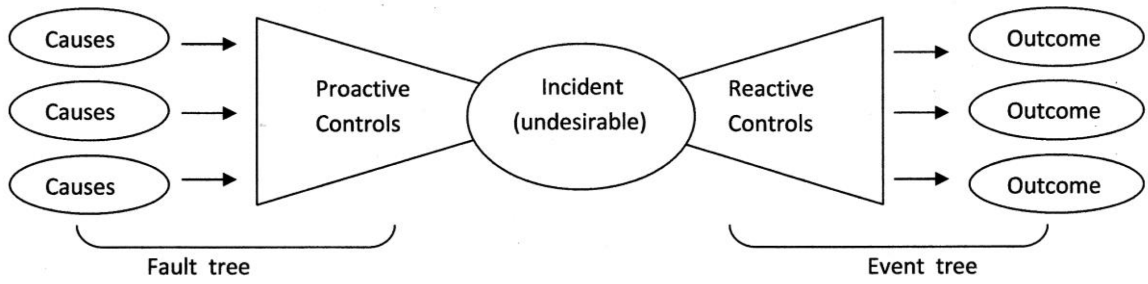

The undesirable event/incident is in focus at the centre of the bow-tie. Proactive controls may be thought of for preventing the possible causes of the undesirable event and reactive controls may be thought of for reacting on the event if it happens to prevent undesirable outcomes. Here a popular tool for analysing causes (fault tree) and a popular tool for analysing consequences (event tree) is indicated as well.

The emphasis of a risk analysis may in some cases be limited to the proactive or the reactive side, and will also depend largely on the field in question and the context. The chosen analytic methods must be adapted to the objective of the study. It is important to understand the risks, and readymade schemes for analysis may not be advisable. Be aware that a proper risk analysis is more than predictions and probabilities. We have to understand the assumptions underlying our models and quantifications. More mathematics and statistics may not always lead to better and more reliable analysis. A crude analysis may often be preferable, as it is not dependent on that many model assumptions that may be questioned. Moreover, it may more easily accommodate qualitative aspects, which may be hard to quantify without assumptions that may limit the validity of the analysis. It is important to be aware of both the strong and the weak side of our analysis. The uncertainties about assumptions, models and stipulated quantities (numbers) may be large, and should not be swept under the carpet. Sensitivity analysis, i.e. analyses made under varying assumptions, may come to help, but cannot cope with all uncertainties involved. A risk analysis should not be limited to the systematization of what we know about the risk involved, but also of what we do not know. Many risk analyses are weak on the latter, partly because the analysts may not know how to handle it, and the principal (employer, decision maker, politician) do not request it. They want definite answers!

Risk treatment: The objective of this step is to select one or more options for dealing with the risks required to act upon. The type of action may be one of the following:

- Avoidance - by abandoning the activity
- Reduction - of probabilities and/or consequences
- Transfer - by insurance or other means
- Retention - doing nothing, by choice or default

The choice of action has to balance the cost and effort of implementation by the benefits in the short and long run. Legal, regulatory requirements and social responsibility may often override any cost benefit analysis in economic terms.

The residual risk after treatment should be reviewed and documented, and where appropriate, further treated. Clause 5 of the ISO 31000: 2009 document also provides some guidance on the preparation and implementation of risk treatment plans. The risk management process should be recorded and be traceable, so that one can learn from experience and improve the process.

We return to risk treatment in section 2.4.

As mentioned in section 1.6 there is a Norwegian standard for risk assessment NS-5814: 2008 ("Krav til risikovurdering") describing the joint process of planning, risk analysis and risk evaluation.

This standard presents requirements to describe background and objectives for the assessment, and deals with ways to organize the work and the relations to different parties of interest. Important elements are the choice of risk analysis method and the choice of data adapted to the context and objective. These choices and simplifying assumptions made have to be realistic, and documented being so. The clauses for risk analysis have requirements for the identification of hazards and undesirable incidents, leading to a risk description to be used in the risk evaluation. This involves requirements for cause-effect analysis and probability assessments. New to the standard is the risk evaluation step involving (i) comparison of identified risks with criteria for acceptable risk (ii) identification of risk reducing measures and their expected effect (iii) conclusions and documentation of the work. The standard states requirements for each of these elements. The conclusions should be precise, unambiguous and robust, and suited for the risk treatment to follow. The standard describes minimum requirements for the written documentation, which also have to make the actual assessment process visible.   The steps covered by the standard are the planning, analysis and evaluation as shown in the graph:

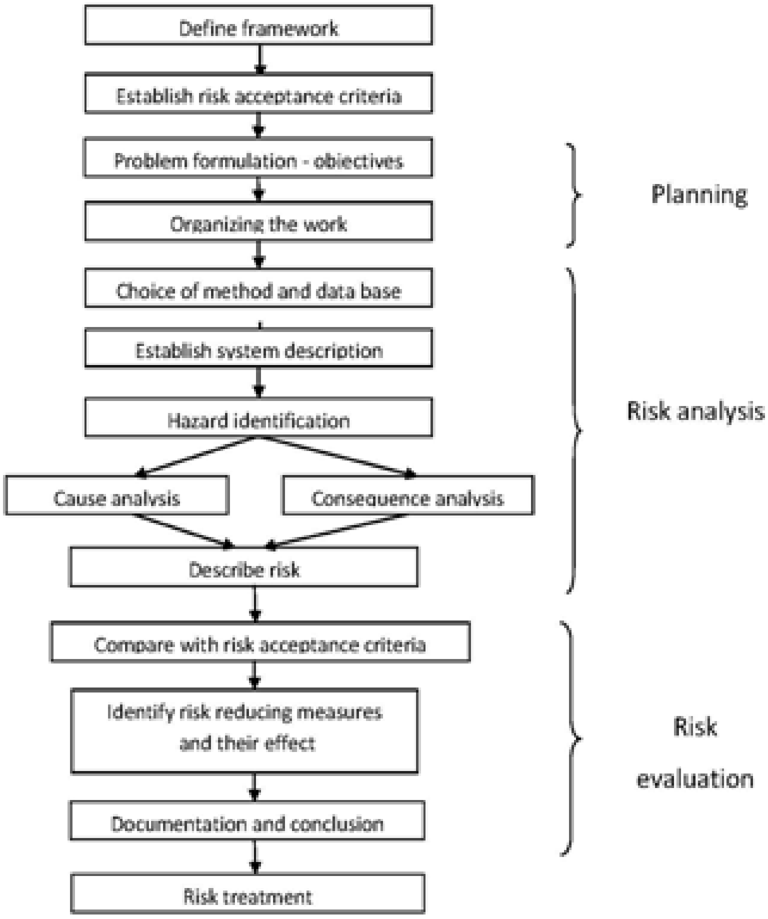
Define framework
Establish risk acceptance criteria
Problem formulation - objectives
Organizing the work
Choice of method and data base
Establish system description
Hazard identification
Cause analysis
Consequence analysis
Describe risk
Compare with risk acceptance criteria
Identify risk reducing measures and their effect
Documentation and conclusion
Risk treatment
Planning
Risk analysis
Risk evaluation

2.2 Risk assessment: Methods and tools

A follow up to the standard ISO 31000:2009 *Risk management – Principles and guidelines* is *ISO/IEC 31010 Risk management - Risk assessment techniques* a standard worked out by International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC). The three main chapters of the document are: Risk assessment concepts, Risk assessment process and Selection of risk assessment techniques. Then follow annexes of informative nature with some details on a variety of analytic techniques.

The context will largely influence the choice of analytic approach for risk assessment: simplified, standard or advanced. The choice among standard methods will also depend on the issue under study and the aim of the study for the stakeholders.

Example: Road tunnel risks
There are many different contexts e.g. combinations of the following: (i) type of tunnel (surface or underwater) (ii) type of tunnel (single or double tube) (iii) sloping (iv) length (v) type of traffic (common or some specific) (vi) traffic loads (sporadic or dense). The context may also depend on whether it is a tunnel in planning, improvement or handling an occurred problem. Planning a safe tunnel, fulfilling the needs of the stakeholders, is entirely different from deciding whether the current fire ventilation system of a specific tunnel is sufficiently dimensioned, or should be replaced. In the former context the issue is mostly the prevention of unwanted events, balancing risks, opportunities and costs. In the latter context it is not so much the prevention and analysis of causes, but the handling of unwanted events, if and when they occur. In some contexts a crude analysis may be sufficient. Other contexts may need a standard analysis, perhaps supplemented by some more advanced methods.

We will here briefly mention some widely used analytic methods.

Analytic methods

The analytic methods for risk analysis may be characterized according to several dimensions in emphasis: Limited issue or integrated issues, stage in the analytic process, crude or detailed, preliminary or not, qualitative vs. quantitative. We may group the analytic methods in three categories

- Schematic methods (FMEA, HAZOP, SWIFT etc.)
- Semi-quantitative visual methods (Event trees, fault trees, Bayesian nets etc.)
- Quantitative model-based methods (Reliability theory, Extreme value theory etc.)

The simplified approach to risk analysis does not go beyond the schematic methods. These methods may also come to use in standard and advanced approaches to risk analysis, then most often at the preliminary stages of the analysis, and is followed up by more quantitative oriented methods.

Schematic methods

Several schematic methods exist. Common to them is systematic use of schemes similar to this:

Element	Unwanted event	Cause	Consequence	Frequency	Risk	Action
: :	: :	: :	: :	: :	: :	: :

Some schemes address all issues from possible unwanted event to possible actions, like this one. Other schemes may focus on specific issues. Some schemes are crude and qualitative, but often sufficient for the purpose. Some are typically used in the early stages of the analytic process, to be supplied by more detailing and quantitative elements later, often by means of other methods. The various schematic techniques are typically used in a group setting, with people with different background. In an industrial context, this may for instance be from design, operation, maintenance and safety. The techniques are not hard to grasp, and they are designed to obtain involvement and promote creativity, in order not to overlook anything. In some cases detailing and quantification are left to trained risk analysts.

Example: Job analysis
Depending on the context, the elements may be characterization of job

- by part jobs in time
- by risk locations for the job
- by the different items handled

Example: Road tunnel risks
Consider just a single element, a specific road tunnel. Unwanted events may be: Motor stop, Car crash, Dropped cargo, Leakage, Overturn, Fire, Drops from ceiling etc. For each unwanted event we may also have more detailed descriptions. Note that here an event may also occur as possible cause of another, e.g. car crash and leakage may be part causes of a fire. How the analysis will proceed will depend on the context established for the risk analysis, e.g. how to dimension the ventilation system, or an even more specific context, whether the current fire ventilation system is sufficiently dimensioned, or should be replaced. In the latter context, it may be found that fire is the only unwanted event to focus on, since the dimensioning of the fire ventilation system cannot affect any of the other events. However, in the initial phase of description and analysis there should not be any restrictions on the creativity, and some suggestions may typically be moved to possible cause for the relevant unwanted event. However, the cause of fire may not be a major issue for the context of dimensioning the fire ventilation system, but may be so in relation to other contexts focusing on risk prevention.

We will briefly mention three schematic techniques:

FMEA – Failure Mode and Effect Analysis is a type of systematic analysis for uncovering possible errors, their consequences and preventive action for components in a system
Some types of schematic analysis resembling FMEA make use of predefined checklists of questions to be asked. Among these are HAZOP and SWIFT, mostly used within the specific areas they were developed.

HAZOP – Hazard and Operability study is a systematic qualitative approach for uncovering weaknesses and dangers in complex systems and processes, where combinations of (minor) unwanted events may have a detriment consequence. Asking the specific HAZOP questions may prevent overlooking such circumstances.

SWIFT – Structured What-IF Technique is an approach for identifying departures from normal operating conditions by questioning What-If in a systematic manner. This tool is flexible and may easily be adapted to different fields and application areas.

Risk Matrix is a way of presenting the hazard level as a product of consequence and likelihood.

In principle, it may look like the following:

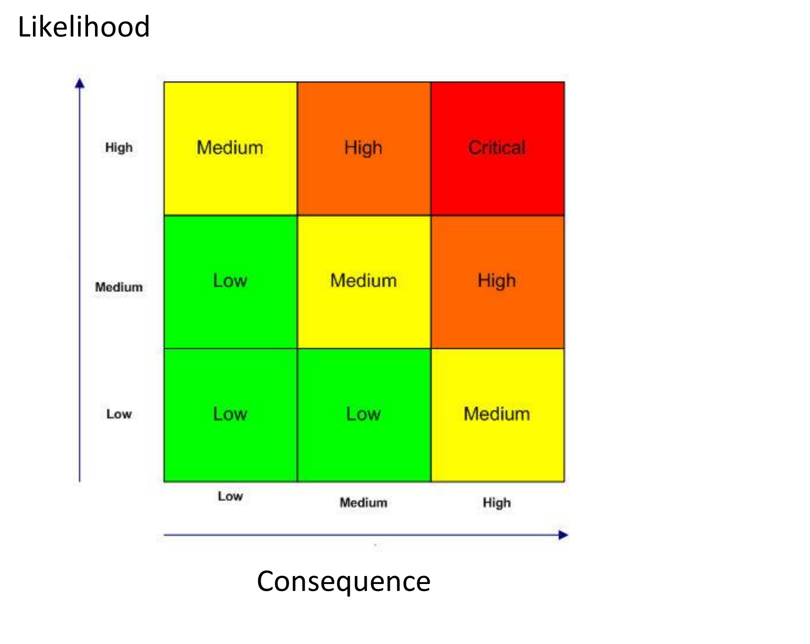

Here three categories (Low, Medium, High) of increasing consequence are along the horizontal axis and increasing likelihood categories (Low, Medium, High) on the vertical axis. Increasing hazard threat then corresponds to going in the North-East direction, and the boxes are given names (Low, Medium, High, Critical) and colored accordingly. The coloring may correspond to the extent of necessary measures to deal with the hazard (risk treatment). In practice the number of categories on the axis is chosen to suit the context, and so is the definition and naming of categories.

Example:
Negligible: One Minor Injury
Marginal: One Severe Injury or Multiple Minor Injuries
Critical: One Death or Multiple Severe Injuries
Catastrophic: Multiple Deaths

Unlikely - Possible - Likely

In some fields, there may be broad consensus on the definition on categories.

A tool of this kind is useful both in simple risk analysis (like the ROS-analysis of section 1.9) and in some more extensive analysis. It may also be useful for presenting summary results to the public, but can be misused.

Semi quantitative visual methods

<u>Event tree</u>

An event tree exhibits the possible events that may follow an initiating event, i.e. the rightmost part of the Bow-tie diagram presented earlier. It may be used for the (i) creation of ideas on what may possibly happen (ii) study and analysis (iii) documentation of the risk picture. Here is a simple example where the initiating event is a fire (B). To stop the fire a pump with water supply is available, as well as a back-up pump. The following events are whether the main pump works (H), and if not, whether the back-up pump works (R). In both cases there is the possibility that the fire is extinguished (S) or not. Each path from the root of the tree (left) to the tip of the branches (right) represents a possible chain of events.

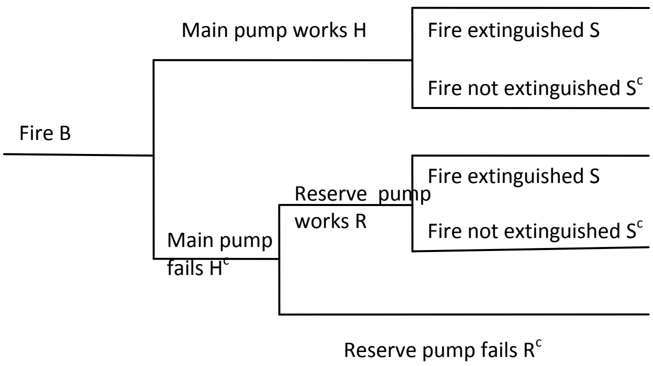

For some purposes just a simple tree is sufficient. For other it is worthwhile to go into more detail. A tree may be limited to a description of the system as it is today, or it may exhibit possible risk reducing measures, e.g. barriers, in order to understand how the system alternatively may work.

The event three can be used as it is, as a graphical representation of the chain of possible events, or as basis for quantification. We then assign probabilities at each branching, representing the conditional probability of each branch, given the events up to the branching point. The probability of each possible chain of events or then given by multiplying the probabilities along the branches

Example: $P(S|B) = P(H|B) \cdot P(S|B,H) + P(H^c|B) \cdot P(R| B,H^c) \cdot P(S|B,H^c,R)$
With P(B) as probability of the initiating event, we have $P(B \text{ and } S)= P(B) \cdot P(S|B)$

Each tip of the tree corresponds to a possible consequence and gets a probability assigned to it. If the consequences are assigned numerical values, monetary or on another common numerical scale (e.g. number of fatalities), we may compute the expected value by the weighed sum of the consequences using the probabilities as weights.

In practice, it may be a challenge to imagine everything that may go wrong, and be able to represent the conditional probabilities fairly, taking possible interrelationships into account.

Fault tree

A fault tree tries to represent graphically the necessary condition for the initiating event to happen, i.e. the left side of the Bow-Tie chart presented earlier. Here is a possible fault tree for the case of fire (B) as initiating event. This is taken as top event of a "hanging tree", and down below are the condition for fire to occur, here leakage (C_1) and ignition (C_2) joined by an AND gate, telling that both have to happen for the top event to happen. To have ignition, we imagine two possibilities: Electronic spark (D_1) or operator smoking (D_2) connected by an OR gate. The bottom event are called the basic events, here are three basic events. The visual difference between the two types of gates should be apparent from the graph. In practice, a number of other standardized symbols are used in fault trees.

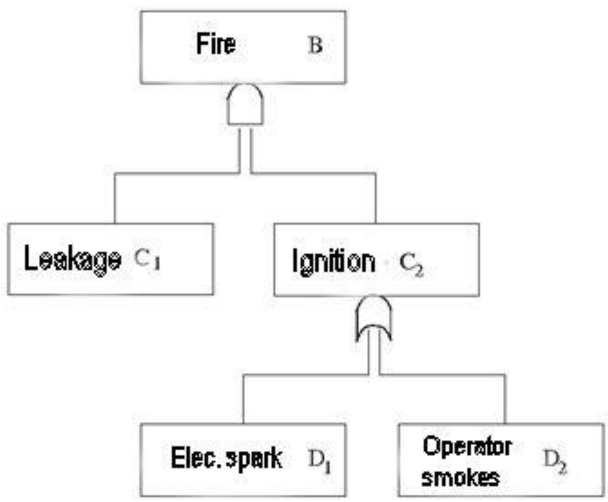

A fault tree may be quite simple or more extensive. We can easily go one more layer down, and ask for the circumstances for ignition, How far to go, depends on the context and aim for the analysis.

The fault tree may be used graphically as it is, or we can use it as basis for quantitative analysis by assigning probabilities to the events. We note that

$$B=C_1 \cap (D_1 \cup D_2)$$

If we assume that Leakage and Ignition are independent events and Electronic spark and Operator smoking are disjoint events, we have

$$P(B) = P(C_1)\cdot P(C_2) = P(C_1)\cdot(P(D_1)+P(D_2))$$

In practice, the challenge is to imagine all the conditions leading to the top event, and assign fair probabilities and make fair assumptions. Here the assumption of independence is vital. Dependencies occur quite often in practice, and when disregarded we may arrive at risk estimates that are far too low.

A fault tree can alternatively be represented by a so-called reliability block diagram. Here follows such a diagram for the situation above, where each box is a possible block for preventing the fault situation (B) to happen. In order for B not to happen both paths from left to right leading to B must be blocked. This happen when C_1 is blocked and at least one of D_1 and D_2 are blocked.

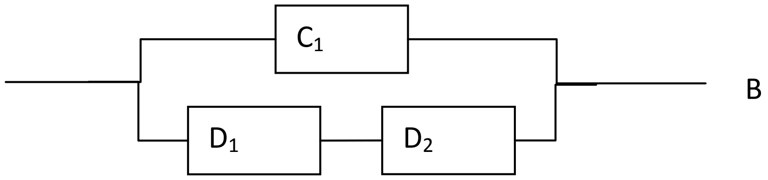

Note that the AND gate in the fault tree corresponds to boxes in parallel and the OR gate corresponds to boxes in series in a reliability block diagram. This is so since it is becomes a reliability diagram for the complementary events to the ones named.

Cause-Effect diagram

Useful tools may also be found in the quality improvement literature, among them are Cause-Effect diagram, used for brainstorming (on the board) possible causes to an undesirable event. Several types exists, the most common is the so-called "Fishbone diagram". Here is a rudimentary one for the event of "Rooms not ready" at a hotel.

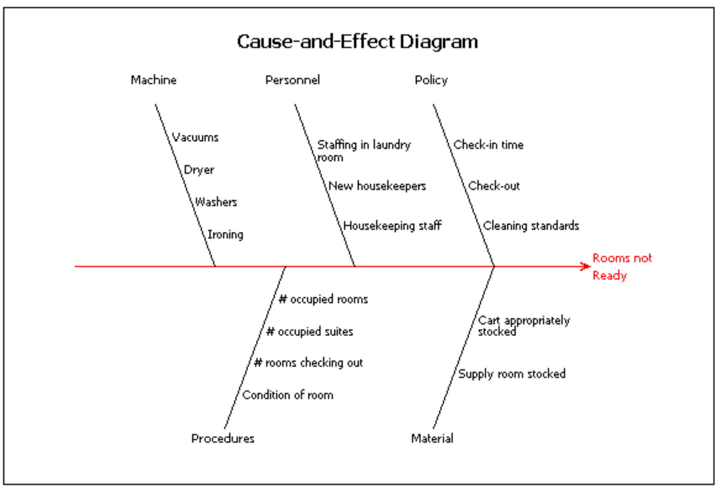

When brainstorming for causes it is often useful to think in terms of the 6 M's: Man, Material, Machine (equipment), Method, Measurement and Mileu (environment). The Cause-Effect diagram has several aims: Create a common understanding of the issue, lay the basis for prioritizing and measurement of characteristics related to possible causes.

Exercise: Tea cooking and serving
A foreign businessman is invited to your home to prepare an important business deal for your company. Well in advance, you start thinking about what may go wrong. Of some importance will be to serve your guest a good cup of tea. Of course, this activity does not deserve more than a sketchy analysis. However, having the upcoming risk analysis training in the company in mind, you decide to make a more comprehensive risk analysis, one that illustrates a systematic approach using simple analytical tools, mostly graphical.

Reliability analysis

A system of n components in series will work if all components work

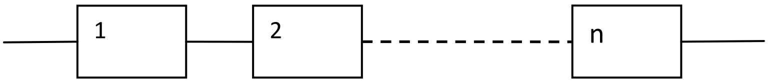

A system of n components in parallel will work if at least one of the n components work

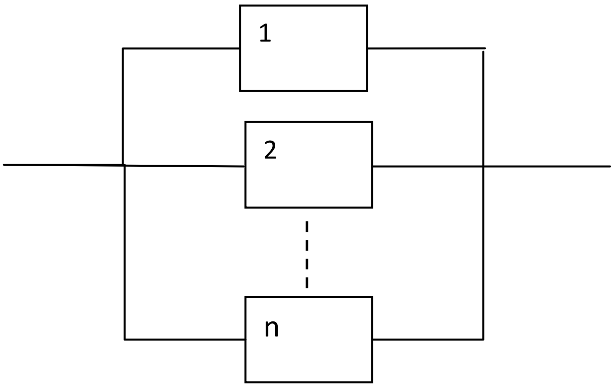

Many systems can be described by a combination of serial and parallel components, where the serial components are necessary for the system to work and the parallel components may be alternative ways of doing things or back-up solutions, say by duplicating components.

In the case that each component has only two states, functioning or not, we can represent this by

$$\begin{aligned} x_i &= 1 \quad \text{if component no. i works} \\ &= 0 \quad \text{if component no. i does not work} \end{aligned}$$

The state of the system may then be described by a vector $\mathbf{x} = (x_1, x_2, ..., x_n)$ of zeros and ones. We may then define the corresponding for the system

$$\begin{aligned} \varphi(\mathbf{x}) &= 1 \quad \text{if the system works} \\ &= 0 \quad \text{system does not work} \end{aligned}$$

We then have

Series system: $\varphi(\mathbf{x}) = x_1 \cdot x_2 \cdot \ldots \cdot x_n = \min(x_i)$

Parallel system: $\varphi(\mathbf{x}) = 1-(1-x_1)\cdot(1-x_2)\cdot\ldots\cdot(1-x_n) = \max(x_i)$

Let $p_i = P(x_i = 1)$. We then have under the assumption of independence

Series system: $P(\varphi(\mathbf{x}) = 1) = p_1 \cdot p_2 \cdot \ldots \cdot p_n$

Parallel system: $P(\varphi(\mathbf{x}) = 1) = 1-(1-p_1)\cdot(1-p_2)\cdot\ldots\cdot(1-p_n)$

A simple example of a system with both components in series and parallel is

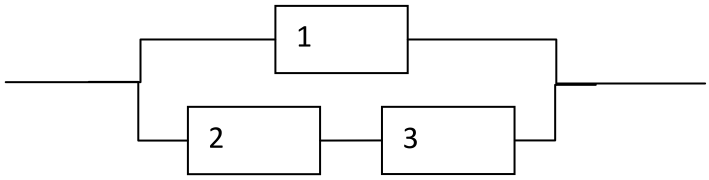

Here $\varphi(\mathbf{x}) = 1-(1-x_1)\cdot(1-x_2 \cdot x_3)$ and $P(\varphi(\mathbf{x}) = 1) = 1-(1-p_1)\cdot(1-p_2 \cdot p_3)$ for independent components.

In practice components do not necessarily fail independent of each other, you may have both positive and negative dependencies. Positive dependency may occur due to a common cause or due to a domino effect. Negative dependency may occur when some failure upstream reduces the load downstream. The modelling of dependencies is challenging, and various approaches are given in the literature.

The discussion above is limited to a static view, and can be extended to the time context, where the life time of each component is modelled by some stochastic process, the simplest one being the Poisson process. This rapidly gets complicated, and we often have to resort to simulations instead of analytical studies.

Statistical methods
A wide range of statistical methods may come to use for risk analysis, in fact most of what you have learned in elementary and intermediate courses in statistics, and many advanced methods as well, depending on the context. Among the methods of main interest are: Multiple linear regression, analysis of variance, categorical regression (logit, probit), time series analysis, panel data analysis, statistical process control, sampling inspection and epidemiological methods. Of potential interest are also have theories linked to statistics, like stochastic process theory, reliability theory, extreme value theory etc.

Example: Distributions
Consider the challenges of risk analysis of socio-technical systems briefly described in section 1.10. The occurrence of outage or breakdown may typically happen at random instants like a Poisson process at a rate determined by data, if available. Sometimes this is taken as an assumption, and with scarce data some expert judgment may help to fix unknown parameters to

obtain a completely specified model. The assumption that the rate is constant over time may not be always justified. Outages due to adverse weather conditions (storms or lightning) may be more common in some parts of the year than others. For the time to restoration X after a breakdown (downtime), we may ask for its (cumulative) probability distribution $F(x) = P(X \leq x)$, or in the current context rather the complementary (survival-) function

$$Q(x) = P(X > x) = 1 - F(x)\,.$$

In practice, it seems that many systems have frequent short downtimes and less frequent long downtimes. This means that the downtime distribution is skewed with long right tail, far from being normal. Distributions of this kind are the Gamma, the Weibull and the lognormal distribution. Of interest are also the Pareto distribution and the so-called extreme value distributions. In our context, we are typically interested in the upper tail of the distribution beyond some (extreme) level q, i.e. $Q(x)$ for $x > q$. The heaviness of the extreme tail of the distribution may then be crucial. Distributions with heavier tails than the normal may have so-called exponential tail (like the Gamma distribution) or even more heavy with power tail (like the Pareto distribution). This means that $Q(x) \sim k\mathrm{x}^{-a}$ for large x, in the sense that the ratio tends to one as x tends to infinity. Here the parameter α express the heaviness of the tail. Software exists for assessing how well data fit to specific distributions, and may also direct you towards a good choice of model.

<u>Example</u>: Explanatory modeling
We want to investigate how some combinations of the variables in $X = (X_1, X_2, \ldots, X_r)$ may trigger, predict or explain an adverse event, here denoted by Y=1 if the event occurs and 0 otherwise. For risk management in an enterprise, we may think of two different contexts:

In system planning: We try out different combinations of input variables, and observe whether they lead to a predefined adverse event or not.

In investigation: An adverse event of a given kind has occurred repeatedly, and we pick up the background information for each event, and at the same time, we collect sufficient background information for situations not leading to the adverse event.

A common model is the logistic model: For given $X = x$ where $x = (x_1, x_2, \ldots, x_r)$

$$P(Y = 1 \mid X = x) = \frac{e^z}{1 + e^z}$$

where $z = \beta_0 + \beta_1 \cdot x_1 + \beta_2 \cdot x_2 + \ldots + \beta_r \cdot x_r$ is a linear score-function to be estimated from data.

Data mining methods

These methods are of exploratory nature, and are typically designed for situations with large amounts of data, both with respect to the number of variables and the number of cases observed. Thus they are often referred to as computer intensive methods. Among these are Classification and regression trees (CART), Bayesian Belief Nets and Neural Nets. Such methods have a potential in risk management as well (or claim to have).

Example: CART

Consider breakdown (Y=1) for given background information X=(X_1, X_2,..., X_r), where the background variables may be of any kind: categorical (nominal or ordinal) and/or continuous. We want to find the variables that may trigger breakdown, individually or in combination. In the case of many variables and possible non-linear relations and interaction (and no theory) this is a difficult task. Classification and regression trees are binary splits according to an appropriate criterion for variable selection. As an illustration the output of such a procedure we give the following graph:

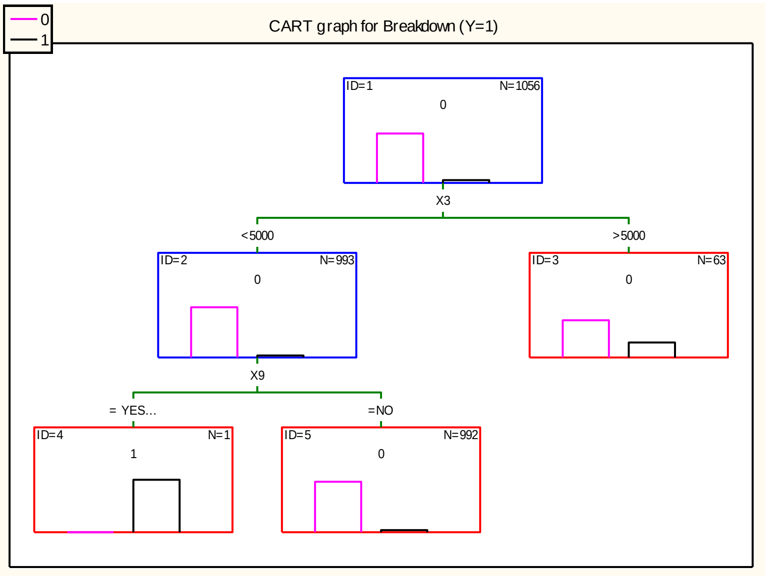

Here we had r=18 variables and observed N=1056 cases, 58 of them were breakdowns. The first split was with respect to the continuous variable X3, where large values seem to be the risky ones. However, for small values, watch out when the categorical variable X9=YES.

Exercise: Continue the interpretation of the tree in this example.

A wide range of method and tools are marketed as useful for risk analysis. However, some of them beyond reach of non-experts. Before we leave this section of analytic methods it is appropriate to keep in mind:

- Do not make risk management too advanced!
- Simplistic methods may be sufficient and recommended in many contexts!

Efforts have been made to establish risk assessment as a science on its own right, founded on its own paradigm and with its own quality criteria. According to Aven (2011)[9] the science of risk assessment is “the development of concepts, principles, methods and models to identify, analyze and evaluate (assess) risk in a decision-making context”. As in many other sciences the scientific methods of risk assessment have to satisfy some criteria for reliability (consistency in use) and reliability (address and answer the relevant question). There are several ways of clarifying these concepts in order to make them relevant for a risk management context. In doing this one have to have in mind the main objectives of a risk assessment, which may depend on the context:

- To provide an “objective” knowledge description of unknown quantities.
- To provide a faithful representation of uncertainties based on available information and knowledge.
- To provide expert judgments about unknown quantities.

A widely accepted scientific framework must encompass these different objectives, and preferably, also accept different views on the use of probabilities, among others that some uncertainty issues are not easily expressed by probabilities alone. A logical consistent framework encompassing these objectives is given by Aven (2011), where also different views on these issues are thoroughly discussed. In brief, his framework is as follows: A risk is described by (B, C, U, P, K), involving the possible adverse event B, the consequences C, knowledge based (subjective) probabilities P, uncertainties not captured by P, and the background knowledge which P and U are based on. Here the P and U may involve models of some kind, based on the background knowledge. In this set-up probabilities are not taken as the risk per se, but regarded as a tool in the risk description. The framework is essentially Bayesian, and objective probabilities are only used when they are justified, e.g. by symmetry or exchangeability arguments.

The annexes of the *ISO/IEC 31010: Risk management – Risk assessment techniques* provides some informative details on 31 analytic techniques, each presented in the following format: Overview, Use, Inputs, Process, Outputs and Strengths and limitations. The techniques presented are

B1 Brainstorming
B2 Structured or semi-structured interviews
B3 Delphi
B4 Check-lists
B5 Primary hazard analysis
B6 Hazard and operability studies (HAZOP)
B7 Hazard Analysis and Critical Control Points (HACCP)
B8 Environmental risk assessment
B9 Structure « What if? » (SWIFT)
B10 Scenario analysis
B11 Business impact analysis
B12 Root cause analysis
B13 Failure mode effect analysis
B14 Fault tree analysis
B15 Event tree analysis

B16 Cause and consequence analysis
B17 Cause-and-effect analysis
B18 Layer protection analysis (LOPA)
B19 Decision tree
B20 Human reliability analysis
B21 Bow tie analysis
B22 Reliability centred maintenance
B23 Sneak circuit analysis
B24 Markov analysis
B25 Monte Carlo simulation
B26 Bayesian statistics and Bayes Nets
B27 FN curves
B28 Risk indices
B29 Consequence/probability matrix
B30 Cost/benefit analysis
B31 Multi-criteria decision analysis (MCDA)

[9] Aven: Quantitative Risk Assessment: The Scientific Platform. Cambridge University Press.

2.3 Risk description and modelling

The aim of a quantitative risk analysis is to arrive at some risk description for a system. This will depend on the analyst's background knowledge, experience and available data. We will suppose that the risk description includes the prediction and uncertainty assessment of an observable system performance variable Y of prime interest, which may depend on one or more observable system variables represented by a vector **X**. The framework for establishing the risk description is (see Aven, 2001):

- Observable variables Y (primary) and **X** (secondary)
- A deterministic relationship between Y and **X**: Y=g(**X**)
- Uncertainty assessments of Y and **X**

This is illustrated graphically as follows:

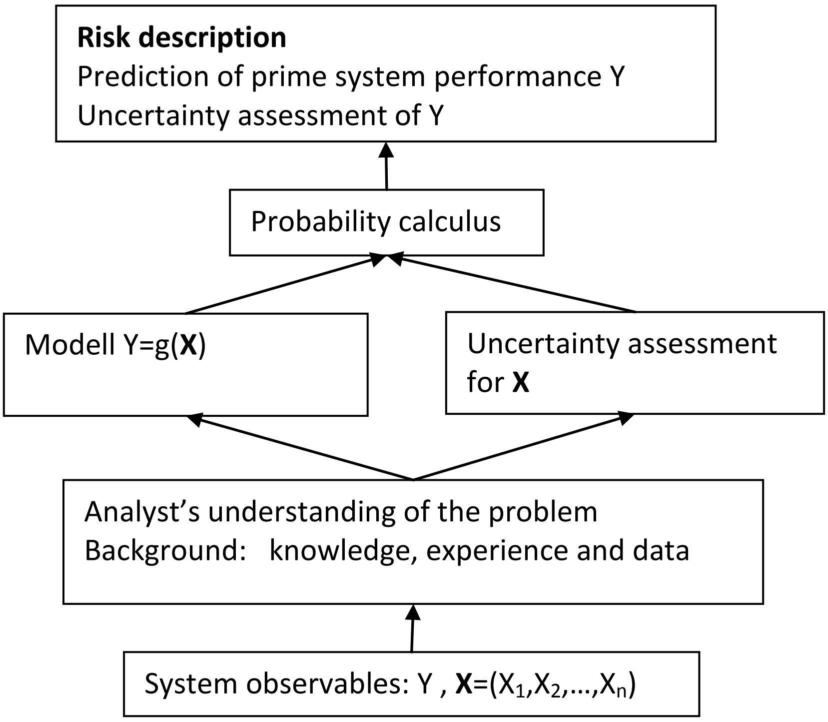

The uncertainties involved should be taken as the predictive uncertainty, taking into account the background knowledge.

What is an observable variable? The notion extends to potential observable. Many risk issues relates to fractions and averages. Are they observable? It depends! The mean of a key quality characteristic of units or the fraction below a given quality level produced on a specific machine qualifies for being observable. However, such a mean or fraction for all similar machines out there is a construct in our minds, and does not qualify to be observable.

Direct assessment of Y

An assessment of the uncertainty in the outcome for an observable Y is by specifying a (cumulative) distribution $F(y) = P(Y \le y)$. This assessment will be based on some background information, which may be part data $y_1, y_2, ..., y_n$ and part other knowledge.
Different types of assessment are available, and may be used depending on the situation:

1. Hard data with little other knowledge: Assign empirical distribution
2. Little relevant data available, but analyst willing to condense other information and make direct probability statements.
3. Conditional assessment and Bayesian thinking combining data and prior beliefs.

A more elaborate and costly alternative to 2, is expert elicitation requiring a panel of diverse expertise, using formal procedures in an effort to fulfil some principles: like reproducibility, neutrality and fairness. These experts will typically not cover all aspects of the risk situation. They will not necessarily reach consensus, and will not necessarily provide results on a probability scale. It is therefore the risk analyst responsibility to weigh and translate the expert opinions. We will not go into more detail on this.

Let us look at each of the three approaches above.

Method 1: Empirical distribution assignment:
Assuming the data $y_1, y_2, ..., y_n$ to be relevant for future uncertainty of Y, we assign the cumulative probabilities by using directly the empirical cumulative counts, i.e.

$$F(y) = \frac{1}{n}\sum_{i=1}^{n} I(y_i \le y)$$

where *I* is the indicator function, being 1 if the argument is true and 0 otherwise, so that the sum is the number of observations less than or equal to y. The empirical cumulative function will have form as a staircase with steps up at each observed value.

Note that we do not assume that there is a true underlying distribution, and that the data are independent observations drawn from this distribution. We assume no more than the data should be relevant for future uncertainty about the outcome of Y, and then this assignment seem to be the best we can do without further knowledge.

Some special cases:

Y discrete variable: $P(Y = y) = F(y) - F(y^-)$ i.e. the step up at y

Y Bernoulli 0- 1 variable: $P(Y = 1) = \frac{1}{n}\sum_{i=1}^{n} y_i$ i.e. the fraction of ones.

For Y continuous variable it may be reasonable to smooth the discontinuous cumulative distribution function to make it continuous as well. This can be done either by no assumption on distribution class or by some specific assumption, e.g. normal, lognormal, Gamma etc. if prior knowledge supports it.

The number of observations needed for an assignment of probabilities of the kind above, depends on the context. In cases when the extreme of the distribution is of minor importance, as low as 10 may be sufficient. If the extreme of the distribution is the key issue, hundreds and even thousands may be required.

Method 2: Direct probability assignment by analyst

Either: Choose n numbers $y_1 < y_2 < ... < y_n$ and ask for n assignments so that

$$P(Y \le y_1) \le P(Y \le y_2) \le ... \le P(Y \le y_n)$$

or the complementary ones $P(Y > y_1) \ge P(Y > y_2) \ge ... \ge P(Y > y_n)$

Or : Choose quantiles $q_1 > q_2 > ... > q_n$ and ask for $y_1 < y_2 < ... < y_n$ so that $P(Y > y_i) = q_i$

The latter approach may be preferable, and typically only a few assessments are needed, e.g. when extreme risks are of minor importance for 0.90, 0,75, 0.50, 025, 0.10. If the assessment can be combined with specific distributional assumption, we can in principle fit a distribution by taking n as the number of parameter of the distribution class, e.g. n=2 for the normal. However, it may be worthwhile to extend this somewhat, in order to check agreement with the implication of the first choices. Recommendations for carrying out such assessments may be found in the literature.

Method 3: Conditional assessment and Bayesian thinking
In some contexts it may be easier to think about the probabilities related to the observable Y given the outcome of some "state of the world" variable θ . If we can come up with a probability assessment $H(\theta)$ of this variable the consistent assignment of $P(Y \le y)$ would be

$$P(Y \le y) = \int P(Y \le y \mid \theta) dH(\theta)$$

i.e. the conditional assessment is weighed by the distribution H and integrated to give the unconditional assessment. $H(\theta)$ may be referred to as the prior distribution of θ, and if we have data that is dependent on this state of the world we may use Bayes law to update the prior, and thus giving a posterior for θ which can enter the formula above instead. Mathematically it may be convenient to pick probability models that gives posteriors of the same type as the prior and integrates out nicely to give interpretable analytical formulas for $P(Y \le y)$. However, if we do not need that, such problems are easily solved numerically

Example: We want to assess the uncertainty attached to the measurable quality Y of a unit to be produced, and feel that it is critically dependent on whether on the delivery time T of some equipment to be used. Given this delivery time T=t, it is easier to assess $P(Y \le y \mid t)$, and we may be willing to assume the delivery time to be distributed Gamma, and assessed by method 2 above.

Assessment of Y via assessments of X

This will depend heavily on the context and type of model, and is best illustrated by examples.[10]

Example 1

The total cost Y may be expressed as the sum of the costs of n cost components as

$$Y = S_n = X_1 + X_2 + \ldots + X_n$$

Some of these cost components may be dependent, e.g. due to common oil price.

Example 2

Let N_t be the number of customer in the period [0,t] and X_i the quantity sold to customer no.i i=1,2,3,… The total quantity sold in [0,t] is then

$$Y = S_{N_t} = X_1 + X_2 + \ldots + X_{N_t}$$

Example 3

In a service system there is at times no one being served and no one in the queue. Let X_t be 0 if there is no one in the system at time t and 1 otherwise, which is modelled by a given arrival and service structure The fraction of time busy in the period [0,T] is then

$$Y = \overline{X}_T = \frac{1}{T}\int_0^T X_t dt$$

Example 4

If T is a lifetime of a unit, we may think of $F(t) = P(T \le t)$ as an observable quantity, namely the fraction of similar units with lifetime less than or equal to t. In our framework we may take Y=F(t). Suppose we can justify exponential lifetimes, i.e. $F(t) = F(t \mid \lambda) = 1 - e^{-\lambda t}$, where λ may be interpreted as the average number of failures per exposure time, which may be taken as observable and playing the role of X and assigned an uncertainty distribution.

Example 5

The return on a portfolio of n assets, where the fraction of capital w_i is invested in asset no i.

$$Y = w_1 \cdot X_1 + w_2 \cdot X_2 + \ldots + w_n \cdot X_n$$

Here the covariances are crucial for the reduction of risk by diversifying.

In these examples we have a superior observable variable Y expressed as deterministic function of one or more system variables. Based on uncertainty assumptions on these the distribution of Y may be derived.

[10] Some examples may require knowledge of probability distributions beyond the elementary level (see section 3.6).

Example 6

Consider a gas pipe system where leakage (initiating event L) may occur leading to fatalities if the leakage is followed by ignition (A) and explosion (B) and not else, as indicated in the graph.

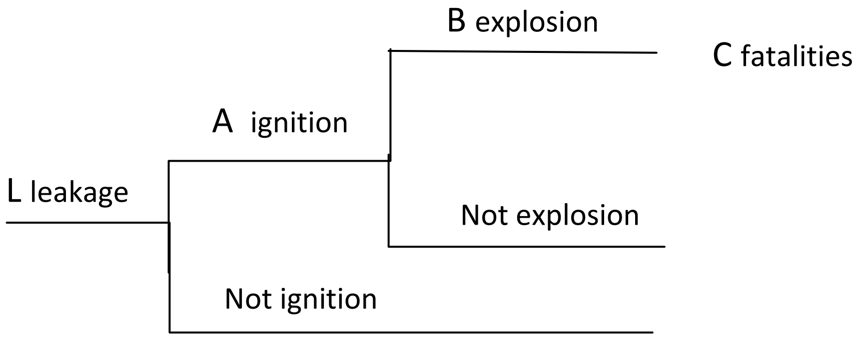

Suppose we have probability assessments for a leakage as follows

$$P(A \mid L) = 0.002 \quad P(B \mid L, A) = 0.1 .$$

and that the expected number of fatalities given leakage, ignition and explosion is

$$E(C \mid L.A, B) = 2$$

With year as time period, there may be several leakages, and we will try to assess the probabilities that relates to the total number of fatalities Y in the coming year. Let N be the number of leakages in the coming year, and suppose that the expected value of N is assessed to be 4, i.e. $E(N) = 4$. The expected number of fatalities in the year will then be

$$\begin{aligned} E(Y) = &\quad E(N) \cdot P(A \mid L) \cdot P(B \mid L, A) \cdot E(C) \\ = &\quad 4 \cdot 0.002 \cdot 0.1 \cdot 2 = 0.0016 \end{aligned}$$

Knowing that the Poisson distribution is used to represent the number of occurrences of random events in a given time span, we may decide to take probabilities from a Poisson distribution with expectation 0.0016. This leads to

$$P(Y = 0) = 0.998401 \quad P(Y = 1) = 0.001597 \quad P(Y = 2) = 0.000001$$

with even more negligible probabilities for more than two fatalities. However, by doing this we have ignored the type of distribution for both N and C. A natural distribution for N may be Poisson, but what about the distribution of C? Suppose the extreme case where the analysis relates to a location where two operators are always on duty, and that if something happens, it will affect both, i.e. C has the sure distribution at 2. Then Y will have possible values 0, 2, 4, 6, …. with probabilities for 2k taken for k for the Poisson distribution with expectation 0.0008 (in fact this is exactly true with the Poisson assumption on N). We then get

$$P(Y = 0) = 0.9992 \quad P(Y = 2) = 0.0008$$

with negligible probabilities for more than two fatalities. The risk of two fatalities is now 800 times the one assessed above. A more realistic assignment number will be somewhere in between these two numbers. In general we have learned that we cannot go by expectation alone and ignore the distribution related to the terminal event.

This example may also illustrate the terminology introduced in this section. The observables Y and **X** are

$$\mathbf{X} = (N, I_{A_i}, I_{B_i}, C_i; i = 1, 2, ...N)$$

where the I 's with subscripts are 0-1 variables indicating whether the subscript event has occurred or not, and the sub-subscript refer to the leakage number. Note that the size of this vector is not known in advance. We may the write

$$Y = \sum_{i=1}^{N} I_{A_i \cap B_i} \cdot C_i = \sum_{\Phi} C_i$$

where Φ is the set of all events where both A and B happen, formally $\Phi = \{i \,;\, I_{A_i \cap B_i} = I_{A_i} \cdot I_{B_i} = 1\}$. By assuming N to be Poisson distributed with expectation λ, it follows that the number of explosions N_e is Poisson distributed with expectation $\lambda \cdot p$, where p is the probability of both A and B happening, i.e. $p = P(A \cap B) = P(A) \cdot P(B \mid A) = 0.002 \cdot 0.1 = 0.0002$, see section 3.10 about thinned Poisson process. Consequently we get $\lambda \cdot p = 4 \cdot 0.0002 = 0.0008$.
Y is a sum of N_e (assumed) independent identically distributed random variables, what is called a compound Poisson process (see section 3.10), for which there is plenty of theory.

Remark. Both the assumption of constant rate of initiating events and the assumption of independent identically distributed C's may be questioned in practice, since an initiating event and an explosion may change the alertness , and some precautionary measures may have been taken.

Exercise
Consider the reliability analysis of a system with two components in parallel, as described in section 2.2. Describe how this fits into the general framework above in each of the two cases:

(i) The primary variable is whether the system works (1) or not (0).
(ii) The primary variable is the probability that the system works.

Note: For risk description in the latter case we need probability distributions over [0, 1].

Among areas of interest and importance not dealt with here are: The modeling of dependent failures and competing risk.

We will return to some selected additional analytic topics in Part 3. Among them are:

- Bayesian methods
- Statistical Process Control
- Sampling inspection
- Extreme value theory
- Risk simulation

2.4 Risk treatment

Risk treatment is the selection and implementation of actions to modify the risk picture according to the risk evaluation. There may be several ways of doing this, most notably

- Risk Avoidance
- Risk Reduction
- Risk Transfer

The left out opportunity of doing nothing is named Risk Retention.

Risk Avoidance
This strategy is the decision not to get involved in or withdraw from a risk situation. In some cases, this means that we have to look for alternatives. Examples are:
At the local level: A decision not to build a gas station next to a kindergarten.
At the regional level: A decision not to build an oil refinery close to bird sanctuary.
At the national level: A decision to abandon all plans for nuclear power.

Risk Reduction
This strategy is to try to reduce of probability of an adverse event occurring and/or reduce the adverse consequences if it should occur. With a risk concept limited to probabilities, the term Risk Optimization may be used instead, where now Risk Reduction in the safety context is about reducing the probabilities.

Risk Transfer
This may be done by insurance, financial contracts, hedging or other means

Risk Retention
This is to accept the loss or gain in the risk situation by choice or default. It also includes the acceptance of risks that have not been identified.

Together these are the risk-based strategies for risk management. We will put this into a classification scheme for strategies in the next section.

An important area for risk reduction is human hazards at the workplace, where the aim is to prevent health risk, injuries and fatalities. Let us look more closely into the case of injury prevention. A common way of thinking about this is the domino theory, by which we imagine that the injury is the last domino piece to fall in a sequence of domino pieces. The logic goes like this:

Effect and Question	Answer
Why do we have an injury? Why did we have an accident? Why did we have mechanical or human error! and so on	We had an accident! An occident occurred from a mechanical or human error! A mechanical or human error occurred as a result of degradation of equipment or ….. and so on

Some say that we should ask Why at least five times, but be aware so that you don't end up blaming society for everything! For this way of thinking we can use the fault tree tool presented earlier.

An alternative way of thinking is the cascade-failure theory, where we start at the other end and ask: What can go wrong at the high level, that affects a multitude of system down below and finally the individuals, not just the single one. An example is failure of electric power supply. For this way of thinking we can use the event tree tool presented earlier. These modes of thinking should be present during risk analysis and risk evaluation as well, but is emphasized here because of their importance also for generating ideas for risk treatment.

In production, construction and material handling, it has turned out useful to imagine hazards as a kind of energy release. This is so in other areas as well, if energy release is given a wide interpretation. Let us see how some risk reduction and prevention strategies may come forward as prevention of energy release:

1. Hazards may be prevented at the product or process design stage.
   Ex. Pressure release valves in gas tanks
   Ex. Non-smoking rule on a ferry
2. Hazards impact may be reduced at the product or process design stage.
   Ex. Limit the attainable power
   Ex. Speed limits on highways, elevated pedestrian crossings
3. Hazards existing after design may be prevented.
   Ex. Limits on allowable filling and pressure
   Ex. Closing the road after heavy snowfall
4. Hazard release can be limited spatially at design stage.
   Ex. Brakes on vehicles
5. The potentially affected can be separated in time and space from hazard at design, and operation.
   Ex. Traffic lights to separate cars and pedestrians
   Ex. Flight corridors and scheduling
   Ex. Use of firewalls
6. The potentially affected can be made more resistant to damage at design and operation stage.
   Ex. Fire and earthquake resistant buildings
7. The damage done may be countered and contained.
   Ex. Sprinkler system and emergency response crew
8. The damaged object/person can be repaired/rehabilitated.

To obtain a reliable system it may be useful to go through the following list of key words, which may provide opportunities: Design, barriers, redundancy, substitution ability, diversity, preventive maintenance, monitoring, procedure reviews, fool-proofing and personnel training.

In the offshore petroleum industry the phrase barriers and barrier analysis is frequently used. Barriers may be designed to

- prevent the occurrence of undesirable events,
- reduce the consequences of an undesirable event.

Barrier design must take into account the operational conditions, and organizational and human factors. Barrier analysis tries to identify (combinations of) risk factors, and study their criticality and for the possible effect of risk reducing measures.

There may be several layers of barriers, in particular for large project, like the North Sea platform operations. The safety systems and other protective measures may be very varied. Here is a general description of layers of different character:

The onion model (from Taggart)

Businesses like the offshore petroleum business face both strict regulations with respect to hazard control and are also dependent on trust among the public, in order to get acceptance from the politicians to move into new and more vulnerable areas, e.g. in the arctic north. The Norwegian offshore accident track record over the last decade is very favorable, with no major uncontrolled blow-outs with severe damage and loss of lives. However, there has been some severe incidents with large damage potential, to be triggered under slightly different conditions. The major Norwegian petroleum company Statoil experienced an uncontrolled gas leak in 2004 on its Snorre A platform, and again in 2010 on its Gullfaks C platform. With this in mind there have been some worries about an insufficient safety culture and inability to learn from experienced near accidents. With this background the Petroleum Safety Authority requested an external assessment of the causes of the last incident. The report (IRIS 2011/156) points to three issues:

1. The tools for safety control have become so complicated that it is threat against safety control by itself.

2. There is a play between several groups, some with no formal authority, over who are in charge of the safety control tools.
3. There is a lack of will or ability to handle criticism within the organization, which may be critical when an incident occurs.

Experience has told that safety is hardly improved by more rules. Although a rational rule-based regime makes the reactions to incidents more predictable and hopefully better coordinated, it may narrow the space for reactions to the unforeseen (cf. the "Iron cage" concept of the sociologist Max Weber).

Question: What can the supervisory authority do to improve safety: "Soft reaction-let them learn" or "punitive reactions-let them pay".

Risk transfer: Insurance

Some risks may be transferred by insurance, which means that we are compensated for the insured adverse event if it happens. An important issue in risk management related to safety is: When to insure, and what is a fair premium?

Let us first look at some of the issues as seen from the insurance provider. The transfer of risk by insurance has five characteristics:

1. Ability to spread risk
   - when premiums from many cover the losses of the few
2. Ability to reduce variance
   - for large portfolio of insured and independent loss events
3. Ability to segregate risk
   - by differing premiums for high and low risk groups
4. Opportunities for risk reduction
   - by bonus system or premium reduction for safety measures
5. Opportunity to monitor and control behavior
   - by inspection and campaigns

In order to transfer risk by insurance the risk event in question has to be insurable. This requires typically that the event must be well defined and estimates of the chances of the event and the size of losses may be established, so that a reasonable premium can be calculated. Insurability does not necessarily mean that coverage can be obtained. The insurance provider may find that the premium necessary to make the product profitable will not attract sufficiently many customers, and then stay out. [11] For some risks, the country itself may take the responsibility, in particular the case of environmental hazards, where means and money for recovery are needed.

Let us briefly look into some aspects of insurance premiums:

The collected premium paid by the insured will exactly balance the payments of the insurance company in the long run, if the net premium equals expected loss on each contract. We then assume that the insurance company carry many risks and operates over a long period of time so that it takes advantage up the "law of large numbers". Implicitly we then also assume that no losses are so big that the insurance company goes bankrupt. The net premium (or pure risk) for damage L in a single contract is therefore given by
P = E(L) which, in order to separate the probability and the consequence, may be written as $P(L \neq 0)E(L|L \neq 0)$. In practice, the net premium is "loaded" to secure that the company stays alive, to cover the costs of doing business and to add a profit to the owners. There are many different principles for loading the net premium, among them to add a factor times the standard deviation of the loss. All of this presupposes that we have available (at least) these characteristics of the loss distribution. Alternatively, the premium may be based on solvency considerations, within a time context. In order to stay solvent and be able to fulfill their obligations to the rest of their customers the total loss L_t in each period t have to stay larger than the reserves R_t and accumulated premiums P_t for that period. This must hold for all periods. It then makes sense to look at the probability of ruin in a specified (long) time period, i.e. $P\{L_t > R_t + P_t \text{ for some } t\}$. Problems like this require stochastic process assumptions. In actuarial science we find numerous models that address this issue, among others the classical Cramer-Lundberg model. This is based on the assumption that the claims instants occur according to a Poisson process, which means that they arrive at random.
More recently, the classical actuarial models have been challenged by insurance economists, advocating that they do not take into account how the companies invest their money.

In practice insurance companies face a number of complications when, selecting, fitting and applying their models, among them ambiguities related to the population at risk, i.e. is it well defined and do we have the basis for calculated the "correct" premium? In some cases the data may be scarce, but even with lots of data there may be problems. Keywords are *adverse selection* and *moral hazard*, which relates to the behavior of those who buy insurance and have bought insurance respectively. Adverse selection means that the buyer knows more about the current risk state than the provider, and moral hazard means that an insured having protection may start to behave more careless. The insurance providers have means to overcome, to some extent, these problems.

An indication of the added premium due to ambiguity of some kind, taking the well specified case as a basis is given in the following table (numbers adapted from Kunreuther et.al.,1997)

[11] This does not mean that someone willing to offer "protection" against one-time events exist, (claiming a high premium), but this is more like betting (e.g. The monster in Loch Ness will be uncovered by 2020).

Probability	Loss	
	Well specified	Ambiguous
Well specified	Life, car driving, fire	Hurricane (+20-30%)
Ambiguous	Rock fall (+20-40%)	Earthquake, tsunami, bioterrorism (+40-80%)

Transferring risk by insurance should not be an easy way out. Every effort should be undertaken to reduce the risk proactively. This way, also cheaper insurance may be obtained, either by documenting these efforts or accepting higher deductibles.

Let us now consider the case of environmental hazards, where big money for rebuilding structures is needed. This issue is becoming more imminent worldwide, as more and more people seem to cluster in hazardous areas. There are striking differences in the opportunities for both government and insurance providers to handle this challenge. In the developed world, insurance is a natural part of risk management in business, and the country itself may have sufficient means to overcome even severe natural disasters. The question is more how to divide responsibilities between national and local authorities or between public and private. In developing countries, the state may not have the means to stage a quick recovery for the harmed people by itself, and there may be no insurance providers willing to offer protection and few able to pay for any. In such cases, international relief organizations and schemes from organizations like the United Nations and the World Bank may come to help. Efforts have been made to come up with different types of insurance schemes to give coverage from natural disasters which, at the same time, encourage risk reducing efforts. Among the possibilities are the creation of regional insurance markets and third party inspections.

2.5 Strategies for risk management

In this section we will present some general strategies and classification schemes that may be helpful for clarification of ideas before decision making, either prior to risk analysis or after risk evaluation and before risk treatment.[12]

Classification of strategies

There are essentially three broad category strategies for risk management:

I. The risk-based strategy : Avoid, reduce, transfer, retention
II. The precautionary strategy: Containment, monitoring, research and development
III. The discoursive strategy: Inform, build confidence, involve the affected

Quite often, the situation requires a combination of these strategies.

Characterization of potential consequences

We may think in terms of five categories of potential consequences

C1. Ubiquity: Does the risk affect few, many or everybody? Geographical dimension
C2. Persistence: Is the risk sporadic or constant, just now or forever? The time dimension.
C3. Reversibility: Can the non-risk situation be restored or not?
C4. Delay: The time between initiating event and consequence.
C5. Mobilization potential: Degree of violation of interests and values of individuals, groups and society.

Characterization of uncertainties of consequences

We may think in terms three dimensions for the uncertainties for the consequences:

u1. Degree of predictability of consequences
u2. Degree of possibility to measure consequences
u3. The persons and groups who assess or perceive the uncertainties

The appropriate risk management strategy, or combination thereof, will depend on how the risk picture may be characterized by the listed factors C1 to C5 and u1 to u3.

We may have different combinations (Low, Medium, High) with respect to the potential consequences and their uncertainties. Let us look at the different combinations in the following table, where we find it convenient to merge the three uncertainty categories for small consequence, and then give the seven categories of interest labels from (1) to (7), with corresponding to approximate increasing risk level. Besides this are columns with additional judgement possibilities to adjust the labelling of a given situation. It may be helpful to have the

[12] See Aven (2003), Klinke & Renn (2002)

above characterizations C1-C5 and u1-u3 in mind when we position a given risk picture in this table.

	Potential consequences	Uncertainties of conseq.	Risk level	Level of authority	Stakeholder implications	Treatment of societal values
(1)	L	L/M/H	Low	Low	Low	Low
(2)	M	L	:	:	:	:
(3)	M	M	:	:	:	:
(4)	M	H	:	:	:	:
(5)	H	L	:	:	:	:
(6)	H	M	:	:	:	:
(7)	H	H	High	High	High	High

From Aven, (2001)

It may be helpful to have in mind the kind of risk strategy is mostly associated with a given risk level. However, this will also depend very much on the area and context for the risk problem. In the case that we try to classify the situation prior to risk analysis, it will be helpful to have ideas on the type of risk analysis mostly relevant, say

- informal qualitative risk analysis
- formal qualitative/semi-quantitative risk analysis
- formal quantitative risk analysis
- quantitative decision analysis

The scheme may also be helpful when deciding whether the situation is within reach of an established risk control system.

Exercise

Classify the following situations with respect to classification category (1) - (7) and determine what kind of risk analysis and give some risk specific measures that may be appropriate, and decide which category I-III they belong.

(a) Individual: Safe driving to work
(b) County: Safe car commuter traffic
(c) Climate change and lost biological diversity
(d) Production safety at oil field
(e) High-voltage lines in an urban neighbourhood
(f) Build oil refinery based on new technology
(g) Cigarette smoking
(h) Building a nuclear facility
(i) Human interference in the ecosystem

Risk management with the public as a major stakeholder is a challenge**,** in particular when the decision maker faces disagreement between experts and when expert opinions collide with public perception. The experts are needed to provide a risk picture, hopefully reducing the risk of biases, anecdotic evidence and false interpretation of data. However, they are not free from biases. On the other hand, we have individuals and organizations who claim to speak for the public. They may add to the risk picture, partly influenced by their own set of values, and they may disagree as well. Some even have their own scientists. Is it possible to imagine a fair knowledge-based process for developing the decision support material for public risk issues according to identified public values?

An effort in this direction, was made by the German Scientific Council on Global Environmental Change (WBGU, 1999, Klinke and Renn, 2002). A classification scheme slightly different from that above is available, more adapted to risks of nature, technological and environmental. This scheme separates the uncertainties in two parts, the one covered by the probability assessments and the one not covered by the probability assessment, i.e. statistical uncertainty, genuine uncertainty, and ignorance. This scheme puts strong emphasis on whether the occurrence probabilities are known, unknown and highly unknown to science, which may be of major importance in dealing with environmental risks. Six categories are the singled out as the interesting combinations, judged by uncertainty, consequence and other criteria. The classes are given names from Greek mythology related to the myth of Prometheus. Prometheus was regarded as friend of the human race, able to transfer strength, ingenuity and ambition and even foresight to the humans. However, he also had some dark sides. The six categories may be placed in a (Consequence, Probability)-map as follows:

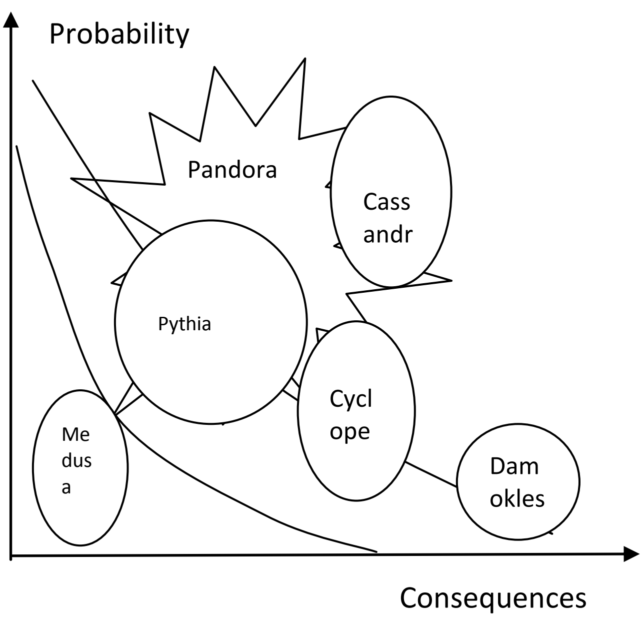

Here curved lines separate three regions:

- The Normal region (Medusa)
  - Low on most risk criteria and common risk balancing methods sufficient

- The Intermediate region (Pythia, but also Cyclope and Damocles)
- The Intolerable region (mainly Cassandra and Pandora)

We see all risk classes except Medusa touches the two outer regions, which means that they pick up aspects that are outside the common simple dimensions for judging risk. The characteristics to define the classes, are given in the next two tables. First, we have a classification with respect to our incertitude: Is the situation characterized by risk/randomness, uncertainty or ignorance?

Type of incertitude	Main criteria	Risk class
Risk/Randomness	Probability of occurrence and extent of damage are *known*	• Sword of Damocles • Cassandra • Medusa
Uncertainty	Probability of occurrence or extent of damage or both are *uncertain*	• Cyclops • Pythia
Ignorance	Probability of occurrence and extent of damage are *highly unknown* to science	• Pandora's box

Then we may separate between classes according to size of probability and consequences or possibly supplementary criteria, as given in the following table, with main group of strategies indicated in the rightmost column (I=risk based, II=precautionary, III=discoursive):

Risk class	Probability	Consequences	Other criteria	Typical examples	
Damocles	Low	High	Not decisive	nuclear energy, dams, large-scale chemical facilities	I
Cyclops	Uncertain	High	Not decisive	nuclear early warning systems, earthquakes, volcanic eruptions, AIDS	I
Pythia	Uncertain	Uncertain	Not decisive	greenhouse effect, BSE-epidemics, genetic engineering	II
Pandora	Uncertain	Uncertain	High persistency	Persistant organic pollutants, endocrine disruptors	II
Cassandra	High	High	High delay	anthropogenic climate change, destabilization of ecosystems, loss of biological diversity	III
Medusa	Low	Low	High mobilization	electromagnetic fields	III

Exercise: Use the descriptions in table and the following hints from Greek mythology to find the probable reason for the naming of the six categories:

Damokles sword:	A symbol of threatening danger in luck
Cyclope:	One-eyed giants, unable to see all
Pythia:	Oracle at Delphi to be asked in case of uncertainty
Pandora's box:	The box with some evils to be kept there
Cassandra:	She gave her prophesy of defeat, but the Trojans did not pay attention
Medusa:	Mythical figure able to menace common people, imaginary but mortal.

Remark. The labelling of strategies may differ between authors. Some use the label cautionary strategy, to single out the risk-based case when some extra protection is in need, to compensate for probabilities not fully known. Some reserve the notion precautionary strategy to the case of scientific ignorance or ambiguity, in line with the strict interpretation of the precautionary principle, like the proactive working definition given about 2004 by The European Environment Agency (EEA) within the European Union (EU).

"*the precautionary principle provides a framework, procedures and policy tools for public policy actions in situations of scientific complexity, uncertainty and ignorance, where there may be a need to act before there is strong proof of harm in order to avoid, or reduce, potentially serious or irreversible threats to health or the environment, using an appropriate level of scientific evidence, and taking into account the likely benefits and drawbacks of action and inaction".*

Here the words uncertainty and ignorance have the meaning given above, while complexity typically means the difficulty to determine causal links among the variables. To this definition, one would like to add ambiguity, meaning differing legitimate interpretations of facts and consequences.

2.6 Non-monetary risks and economics

Many projects involve non-monetary risk and the reduction of human risks may involve monetary costs. A difficult question is then how to weigh non-monetary risks against costs. In this section, we will briefly look into this issue and, among others see the shortcomings of common cost-benefit analysis, and point to some alternatives.

The weighing of cost against risk is illustrated in the following graph, which exhibits three available choices of combinations (risk, cost):

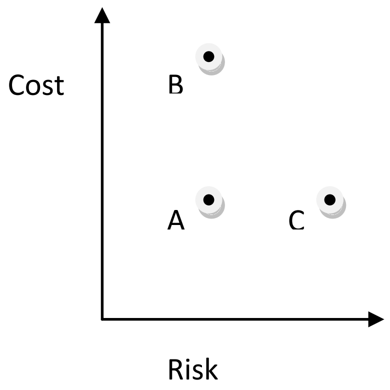

It is clear that A is the optimal choice, but if this turns out not feasible, we are left with the choice between B and C, which is not clear cut, since B has lower risk, but higher cost than C. A risk manager then has to weigh the extra cost against the risk reduction by choosing B instead of C.

Suppose that cost values can be attributed to the risk as well as to controlling the risk. With opportunities on a continuous scale, we can illustrate the situation as in the following graph with two cost curves crossing at an optimal (cost) point.

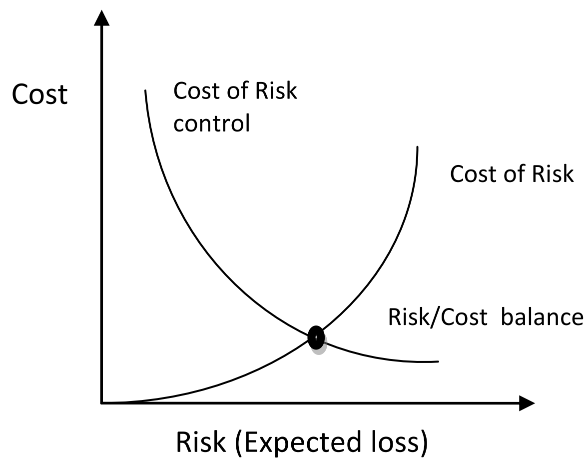

Experts with different background may attack non-monetary risks differently, the main difference being the willingness to translate them into economic terms, in order to integrate all into one or a few key numbers. A safety expert may prefer a multi-attribute analysis including a measure of expected cost per expected saved lives (named cost-effectiveness), while an economist and decision analyst may prefer cost-benefit analysis and expected utility calculations.

Facing non-monetary risks, equipped with tools from economics, questions to be asked are:

- Are expectations sufficient?
- Are net-present value calculations relevant?
- How to do cost-benefit analysis?
- Are utilities practical?
- Is it rational to be risk averse?
- How to face the cautionary and precautionary principles?

Expectations and utility

The use of expected monetary outcome assumes that you are risk neutral, in the sense that the added benefit from an extra dollar payoff is the same whatever the current level of payoff. This is often unrealistic, in particular if the range of possible outcomes is wide. Eventually the value of an extra dollar diminishes. In economics this is connected to the concept of being "risk averse". Economic theory deals with decisions in face of uncertainty, where coherent decisions (in some defined sense) lead to the concept of utility, where the optimal decisions is tantamount to maximizing expected utility (Von Neumann & Morgenstern, 1947). The utility concept represents an effort to put monetary and non-monetary outcomes in a common framework for analysis. Among others, the theory that followed explored the consequences of being risk averse (in some defined sense).

Common utility theory is based on individual preferences, satisfying some axioms of coherent behavior. A consequence is roughly speaking the following: Any outcome A can be represented by a gamble between the best and worst possible outcome, and a certainty equivalent A* in monetary terms can be derived, so that the individual is indifferent between A and A*. Every non-monetary outcome may then be replaced by its safety equivalent, and the optimal decision may be represented by the expectation using the individual subjective probabilities of the outcomes. However, this approach is hardly to any help in our risk management context, at least for two reasons:

(i) A decision theory based individual preference can hardly accommodate the interests of groups, various stakeholders and society.

(ii) It is hard to imagine gambles involving extreme events, and in the case of possible fatalities it sounds unethical to most people.

On the other hand it seems that the insights obtained by economists in terms of utilities and risk aversion are not picked up by the project/safety community.

Cost-benefit analysis

<u>Objective:</u> A basis for choosing among different alternatives and to decide whether or not to initiate an activity.
<u>Feature:</u> Reduce the "value" of each alternative to a single number

– by transforming all "ingredients" to a common monetary scale
– representing the willingness to pay
– taking care of time preferences, and
– aggregating over time

<u>Challenge:</u> Many ingredients are hard to translate to economic terms,
– among them human safety and damage to the environment

Marketed goods are easy to transform to cash, while non-marketed goods require some help from supplementary theory. In particular we face this challenge when we deal with the cost-benefits of projects involving risk of fatalities. In case of possible fatalities it may be possible to get around by use of the concept "value of a statistical life"

Value of a statistical life (VSL) = The expected cost per expected saved life

i.e. the amount of money the Society is willing to pay for reducing the expected number of fatalities by one individual. We may also make use of

Cost-effectiveness index =
Cost added for (life saving) alternative/Expected number of lives saved

<u>Example:</u> Consider a project with the alternative choices A and B each with outcome of type (Benefits mill NOK, #deaths). Suppose that there are just two outcomes for each alternative choice with probabilities given as follows:

Alternatives	A		B	
(Benefits mill NOK, #deaths)	(1,0)	(1,1)	(2,0)	(2,1)
Probabilities	0.99	0.01	0.95	0.05

In case the value of a statistical life is 10 mill NOK we get:

A: $1 - 10 \cdot 0.01 = 0.9$
B: $2 - 10 \cdot 0.05 = 1.5$ (i.e. the preferred alternative)

The cost effectiveness index (for choosing alternative A over B)= 1/(0.05-0.01)=25 (mill NOK/life)
For VSL> 25 mill NOK we choose alternative A.

The value of a statistical life can, in principle, be determined by two different approaches:

- Empirical: Derived implicitly from decisions already made
- Experimental: From designed proposals to reveal willingness to pay

A pragmatic way of including risk into a cost-benefit analysis is by referring to the ALARP principle, and say that something is reasonably practicable unless the costs are grossly disproportionate to its benefits, i.e. when Costs/Benefits > DF, where DF is referred to as an disproportionate factor found reasonable in the area of application. In a case published by the British Health and safety executive (HSE) on the risk of explosion at a chemical plant DF=10 is used, with the requirement that use of a lower DF's has to be justified by the duty holder.

Many projects and investments involve payoffs over time, and some calculations are needed to evaluate the "risk" of the project and evaluate existing alternatives. A common way to perform a cost-benefit analysis, is by discounting future payoff e.g. by the bank interest rate, i.e. compute net present value (NPV).

$$\mathrm{NPV} = \sum_{t=0}^{T} \frac{X_t}{(1+r_t)^t}$$

where

X_t=Cash flow equivalent in period t
r = Discount factor/capital cost/rate of return/time value for money
T=Time horizon

In case of uncertain payoffs we may in principle assign probabilities, and compute expected net present value E(NPV). In practice people often replace the uncertain payoff by their expectation or perceived certainty equivalent (and may call this net present value as well). Note that decisions based on expected net present value neglect the distributional aspects of the cash flow, and thus tell little about the worst cases. A possibility is to do simulations, with draws from the distribution of the X_t 's, which may depend on t, and then establish the distribution of NPV. This gives the opportunity to examine the tail of the distribution as well. In some cases it is natural to take the discount factor and the time horizon to be random as well. Then the calculation of E(NPV) becomes awkward. You cannot just replace them by perceived certainty equivalents. This may largely underestimate the risks. In this case simulations seem inevitable.

Example

A homeowner is living next to a river which is occasionally flooding in the spring, and causes an expected clean-up cost b each year the flood occurs. An investment at cost c is believed to give full protection and eliminate any clean-up. Assuming the probability of flood in any year is p. With a time horizon of h years and discount factor r, the expected net present value of the investment becomes

$$E(NPV) = -c + b \cdot p \sum_{i=1}^{h} (1+r)^{-i}$$

The investment is worthwhile when E(NPV)>0, that is when the benefit-cost ratio >1. We can write this as

$$\frac{B}{C} = \frac{b}{c} \cdot p \cdot d \quad \text{where} \quad d = \sum_{i=1}^{h} (1+r)^{-i}$$

Here follows a table of d for given r and h:

h→	1	2	3	4	5	6	7	8
r=0.05	0.95	1.86	2.72	3.55	4.33	5.08	5.79	6.46
r=0.10	0.91	1.74	2.49	3.12	3.79	4.36	4.86	5.33

If c=10 000 and (b,p)=(100 000, 0.10) we get bp/c=1, and thus B/C>1 for h=2 for both r=0.05 and r=0.10, and are cases where the investment is expected to be profitable with (at least) a two year horizon. For (b,p)= (50 000, 0.02) we still get bp/c=1 and the same conclusion. However, for (b,p)=(20 000, 0.10) we get bp/c=0.2, and the investment seems profitable for a horizon at least 6 years in the case of r=0.05, and a horizon of at least eight years in the case of r=0.10.

Many projects involve outcomes not easily phrased in monetary terms. One may try to overcome this by using a discount factor somewhat higher than the bank interest rate. However, this is somewhat arbitrary, and do not seem very appealing when concerns are about hazards and possible fatalities. Common cost-benefit analysis is based on risk-neutral behaviour, and modifications to reflect risk aversions exist, but are somewhat arbitrary. A possibility is to perform sensitivity analysis, i.e. see to what extent the decision is changed by modified assumptions. When assessing the reliability of a cost-benefit analysis, there are some disagreement about whether to judge “immeasurable risks” (uncertainty about the states of the world) different from measurable risks. For the risk management concept used by Aven op.cit., there is no need for distinction.

Multi attribute analysis

Multi-attribute analysis is a decision support tool to help judging the attributes in a joint context, without bringing these over to a common scale. A simple common form is to use a value function of form

$$v(x_1, x_2, ..., x_n) = w_1 \cdot x_1 + w_2 \cdot x_2 + w_n \cdot x_n$$

where x_i is a measure of the ith attribute and w_i is the weight given to this attribute representing the trade-off between them. Here x_1 may represent some reward in economic terms, x_2 represent a categorical variable, e.g = 1 of no accident occur and = 0 if accident occur. Practical techniques exist to help determine reasonable weights. One can go one step further by defining a “utility function”

$$u(x_1, x_2, ..., x_n) = 1 - \exp(-(w_1 \cdot x_1 + w_2 \cdot x_2 + w_n \cdot x_n) / \theta)$$

where the parameter θ may represent the magnitude of risk aversion, increasing with θ. This may be used in sensitivity studies. The choice of exponential utility is convenient, although somewhat arbitrary, in fact may violate the whole idea of utilitvies representing coherent decisions.

<u>In summary</u>: Common economic analysis based on classic economic theory has limited ability to incorporate extreme events with low probability such as hazards and possible fatalities. It cannot be the sole platform for decision making. A risk analysis based on other principles may be required and added for decision support. The real challenge for the decision maker is therefore how to weigh the decision support material against each other.

2.7 Expert group assessment

In many decision situations involving risk we have little or no “hard data”, and we have to rely on ”soft data” provided by expert opinions. This may typically be the case in project planning and project management of large projects. In some cases the possible outcomes are of extreme nature and have seldom or never been observed before or the project in itself is one of a kind. To get a balanced view you would typically not rely on just one expert, but perhaps a panel of experts with different background. In practice, the experts may have different levels of information sharing. In some cases they may come together and reach consensus In others cases they are separated, and their opinions have to unified by some means. In order to facilitate processes to arrive at a more or less unified judgment, several methods are available. We will in this section briefly expose some group decision method mainly coming from the decision sciences and expert system literature, namely the Nominal Group Technique, the Delphi Method, the Analytical Hierarchy Process and the Decision-Making Trial and Evaluation Laboratory. At the end of this section we discuss some specific issues related to the assessment of probabilities. There is also a literature on expert group assessment in a Bayesian context, where prior beliefs are updated by Bayes’ law. A brief account of this is given in section 3.5.

Before we look into these topics let us mention that data bases of expert opinions exist in some areas of frequent common interest, e.g. on issues related to nuclear safety.

The Nominal Group Technique

The Nominal Group Technique (NGT) is a method that enables a group to generate and rank/prioritize a large number of options, so that consensus is reached in an orderly manner, giving all group members an equal opportunity to contribute. NGT is a good tool for controversial issues or when a group is stuck, e.g. unable to handle many issues. The final result may not be everyone's highest priority, but they will have to live with it.

A successful NGT process requires some preparation. It is lead by a moderator/facilitator who poses questions to each group member, and asking them to prioritize the ideas or suggestions that have come forward from all group members. The process has two parts:

1. Define the issue and generate ideas (collect, clarify and combine)
2. Select ideas: By ranking given by sum of individual rankings (after a sanity check)

It is vital that the process follows strict rules that the participants are familiar with. Detailed prescriptions on how to perform a successful NGT are readily available.

Example: A group of workers has defined the issue as “Why do we produce scrap?”, and have settled on the following voting list

A. Lack of training
B. No documented process
C. Unclear quality standards
D. High labour turnover
E. Insufficient maintenance
F. Lack of cooperation with other departments

The group member are asked to rank to five of them from 5 (highest priority) to 1 (lowest priority). The scores are then collected and summed.

NGT is mainly a small group technique. The advantages of NGT are mainly:

- Generates more ideas than traditional group discussions
- Balances the influence of individuals by limiting the power of opinion makers
- Diminishes competition and pressure to conform, based on status within the group
- Encourages participants to confront issues through constructive problem solving
- Allows the group to prioritize ideas in a democratic manner

However, by its limitation of discussion NGT does not allow for the full development of ideas, and can therefore be a less stimulating group process than other techniques.

The Delphi method

The Delphi method is a structured communication process for a panel of experts administrated by a facilitator. It consists of a number of rounds of information gathering, where each expert answers questions from a questionnaire, which pick up their judgments and their reason for the judgment. This is summarized by the facilitator in a report statistically aggregated results and reasons for judgments, and communicated anonymously to the experts, and is the basis for the next round. This gives the member of the panel the opportunity to revise their judgments. By each round they are likely to get closer to consensus. However, the process is stopped after a pre-defined stop criterion, e.g. number of rounds, achievement of consensus, stability of results. Thus consensus may not be required. The method may be used in small groups as well as in large groups, and they is no need to bring the participants together. In fact, with the use of e-mail and internet the Delphi method may be used to tap the knowledge and "collective intelligence" of very large groups. Some are even talking about the method in the context of a future e-democracy. In order to analyze the panel evaluations the facilitator may use simple or sophisticated statistical methods and modern modes of graphical communication may be used for the back feed of results in each round. Among claimed advantages of the Delphi method are:

- avoids possible negative effects of face-to-face panel discussions
- allows free expression of opinions and encourages open critique
- frees participants (to some extent) from their personal biases,
- reduces the risk of participants sticking to their original beliefs
- reduces the temptation to just follow the leader or the majority
- facilitates admission of errors when revising earlier judgments

The Delphi method has been widely used for prediction based on expert panels, in particular when data is scarce, and we are left with (often self-proclaimed) experts. It has also been used in an early stage in the formation of theories.

The record track of the Delphi method is mixed, which may be due to that future developments often defy prediction by consensus methods. In defense, it is said that in many cases where the method produced poor results, it is just poor administration of the method. It is hard to know. Note also that if panelists are misinformed or ignorant about a topic, the use of Delphi may only add confidence to their ignorance. A drawback with the standard version of the method was its inability to take into account multiple impacts, which typically were considered as being independent. Extensions of the Delphi method exist to remedy this problem, such as *cross impact analysis*. Despite its shortcomings, the Delphi method is a widely accepted forecasting tool and has been successfully applied in many areas. However, the most successful applications so far are forecasting single scalar quantities.

In risk management, the Delphi method may be useful for examining several types of items, not just *forecasting* items but also *issue* items, *goal* items, and *option* items. This may require evaluation on several scales like *desirability*, *feasibility* (technical and political) and *probability*. It is then possible to outline different scenarios: the *desired* scenario (from desirability), the *potential* scenario (from feasibility) and the *expected* scenario (from probability). Many Delphi studies involve the establishment of a tentative list of possible risk factors.

Example (Schmidt et al. 2001)
A Delphi study aimed at a ranked list of common risk factors for software projects as a foundation for theory building about IS project risk management. Participants: Three panels of experienced software project managers from Hong Kong, Finland and the United States. The highest ranked risk factor was "Lack of top management commitment'', where could stimulate the formation of specific research questions related to this particular issue. Another risk factor ''Conflict between user departments,''

The Analytic Hierarchy Process (AHP)
The Analytic Hierarchy Process (AHP) provides a rational framework for structuring a group decision problem, by representing and quantifying its elements, relating those elements to overall goals and evaluating alternative solutions. AHP attempts to mirror the human decision process, and is mostly applied in a group setting. It is widely used, but is still somewhat controversial. The steps of AHP and the steps are as follows:

1. Decompose the decision problem into a hierarchy of more easily comprehended sub-problems, each of which can be analyzed independently. The main objective or decision goal is on the top, and is decomposed into factors or sub-goals and so on, if necessary. The elements of the hierarchy can relate to any relevant aspect of the decision situation, measured or roughly estimated, tangible or intangible, even aspect that are poorly understood. Choices or alternatives for reaching the (sub)goals are added at the bottom, not neccessarily mutually exclusive
2. Identify criteria to evaluate the achievement of objectives/goals. Evaluate the elements of the hierarchy systematically by pairwise comparison of elements that share a common parent in the hierarchy, with respect to their impact on that parent, i.e. importance and strength. In making the comparisons, both hard and soft data about the elements can be used, combined with subjective judgments about the relative meaning and importance of the elements.
3. Convert the pairwise evaluations to numerical values that can be processed to provide overall judgments for the hierarchy (typically on a 1 to 9 scale). A numerical weight or priority is derived for each element of the hierarchy, which is checked for consistency. This allows rational comparisons of very diverse elements in the hierarchy.
4. Calculate numerical priorities for each of the decision alternatives, representing its relative ability to achieve the decision goal,

A hierarchy map with four factors and three choices will look like this.

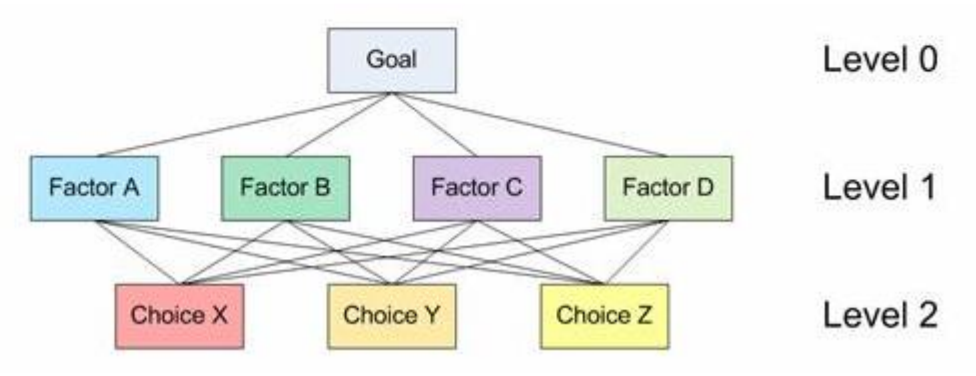

First, we make pairwise comparisons of factors on level 1 with respect to the goal at level 0. With four factors, there will be six such pairwise comparisons.[13] This establishes the priorities of each of the four factors with respect to reaching the ultimate goal. Then we make pairwise comparisons for the three choices of level 2 with respect to each of the four factors of level 1, one at a time. Each of these give weights to the three choices   From these, one compute composite weights for each choice. In some cases the priorities for a factor is so low that it can be removed from the analysis, but then remaining numbers have to be rescaled.

The paired comparisons are often done on a 1 to 9 scale like this:

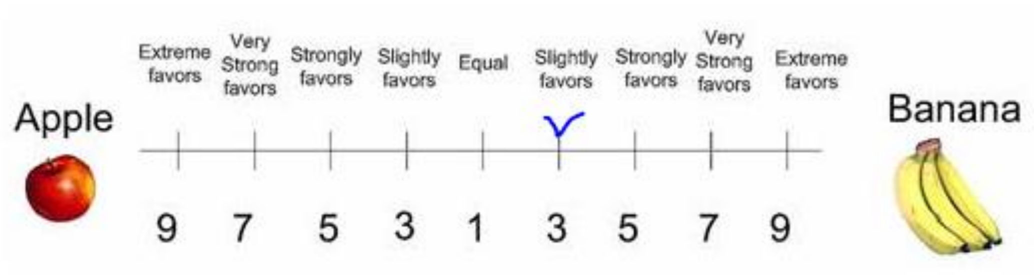

In this case the respondent slightly favored banana to apples, which is recorded as a 3. If the respondent slightly favored apples to bananas that would be recorded as the inverse 1/3. With three choices, say apple, banana and pear, each respondent will have just three comparisons to make, and the result of the these judgments may be recorded in a 3x3 matrix, where we see elements on each side of the diagonal are inverses of each other, and we recognize the 3 as banana over apple, and 1/3 as apple under banana.

$$
\begin{array}{c c}
 & \begin{array}{ccc} apple & banana & pear \end{array} \\
\begin{array}{c} apple \\ banana \\ pear \end{array} & \begin{bmatrix} 1 & 1/3 & 5 \\ 3 & 1 & 7 \\ 1/5 & 1/7 & 1 \end{bmatrix}
\end{array}
\qquad
\mathrm{w} = \begin{bmatrix} 0.2828 \\ 0.6434 \\ 0.0738 \end{bmatrix}
$$

From this the so-called priority vector **w** is calculated by normalizing each column, so that the elements sum to one (i.e. divide each element by the column sum, and then average the three

[13] In general with n choices there will be n(n-1)/2 pairwise comparisons.

columns thus obtained (more sophisticated calculations exist) These numbers are a summary of the facts that bananas are preferred over apples, which in turn is preferred over pears, so the rankings are consistent in this sense. These numbers can now be used as weights in the further analysis.   We can also derive a measure of consistency, which may be needed in situations with many choices, when it is hard to be consistent. This way we may put weights to the different respondents, and also relate them to limits of unacceptable inconsistency.

Example (Söderholm & Nyström, 2009)

In a study made for the Swedish and Norwegian railway administration AHP was tried out on a group of six track managers. The main objevtive was to find the characteristics of a sosio-economically efficient railway system, within the context of prioritizing their maintenance. They came up with 8 criteria: Cost, Safety, track work time, punctuality and availability, condition, environmental impact, own abilities and development, collaboration with stakeholders. The example in the graph below is limited to the first four sub-goals/criteria with 5 different modes of actions.

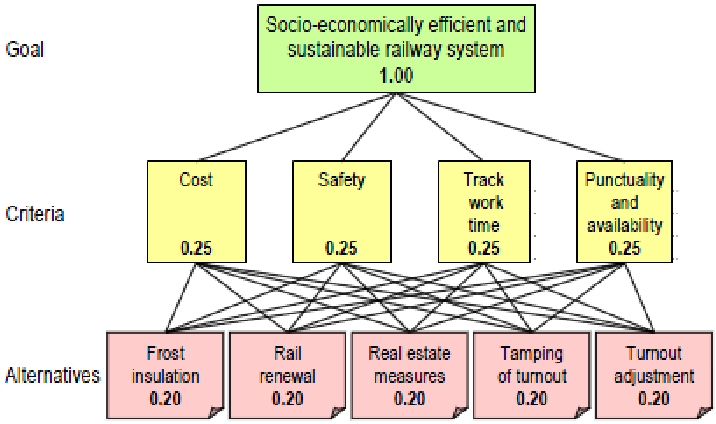

Here the numbers in the boxes are the weights prior to the experts judgments. We will not go into details on what they found, beyomd that Safety turned out to get the highest weigh for the criteria, and there was regarded as the most important one.

Expert judgment of causation - DEMATEL

A risk assessment sometimes requires clarification of cause-effect relationships, in particular for establishing the basis for a successful risk treatment. Data may help to clarify, but data may be scarce. Even with abundant data, it may be difficult to single out the true causes or causes that can be manipulated to reduce risks. As pointed out in most statistical texts: Correlation does not necessarily mean causation! From statistics, we are told that a controlled experiment is the preferred setting for demonstrating causation. However, in most cases, this is not feasible, and we are left with circumstantial data only. Then it does not necessarily help to have lots of data. Statistical science has demonstrated the limitations, but has also come to help in recent years to get a better grip on the possibilities of inferring causation from circumstantial data.
In the case that data is not decisive, scarce or not existing, we have to rely on expert opinions, preferably a panel with varied background and knowledge. The question is then how to uncover and unify the different viewpoints of the members of a panel.

A possibility is the DEMATEL method (short for Decision-Making Trial and Evaluation Laboratory), coming from the expert system literature. It has been around from the early 1970's[14] and applied in many areas, in technology, environmental management and even anthropology. The method allows the determination of the relationships between interacting evaluation criteria and determines a value structure. By this, one can arrive at the factors believed to be must beneficial in order to facilitate change, in our setting risk reduction. The steps of DEMATEL are mainly as follows:

1 The panel starts by coming forward with a number of factors or criteria that may influence each other, including the ones that in focus for successful operation or successful change. Some may represent different aspects of a sub-group or sub-system, say technical, organizational, management, social, cultural, external, cost.
2 Then each member of the panel state their opinion about each pair of factors on a verbal scale, say No influence (N), Very low influence(VL), Low influence (L), High influence (H), Very high influence (VH) .
3 The choices are then converted to numbers in order to integrate the various experiences, opinions, ideas and motivations.

The outcome of the analysis is a grouping of cause criteria and effect criteria, together with measures of the strengths of the relationships. The method also helps to visualize the causal relations among groups of factors by an impact-relation map (IRM). An extension of DEMATEL came forward in the 2000's, where the verbal statements are taken as vague, and the preciseness of statements is represented by fuzzy numbers handled by fuzzy set calculus, Wu & Lee (2007).

Example Airline safety (Liou et.al, 2008)
The following 11 factors were established

Factors	Functions
Communication	Language barrier, crew resource management, maintenance resource management
Documentation	**SOP's, procedures, standards, audit reports, assessment findings, regulatory requirements, incidence registration**
Equipment	**Tools, plants, other required equipment maintenance or calibration**
Incidence investigation	**Contributing factors , human error risk, event and remedy cost, rating alternative remedies, corrective actions**
Safety policy	Setting organizational structure, roles and responsibilities, plans and managers commitment to safety
Rules and regulations	**In-house safety rules and regulationsenforced by administration**
Safety commitment	**Develop strategies for safety, supervisingsafety plans, corrective actions, sub-contractors, and allocating resources for safety improvement**
Safety culture	**Organizational values, beliefs, legends, rituals, mission goals, performance measures, responsibility to employees, customers and community**
Safety risk management	**Internal audit, hazard/risk identification analysis/assessment and control, compliance with legal and other requirements**
Training and competence	**Initial and recurrent training, individual in the positions meet competence requirements**
Work practice	**Flight operations, maintenance, ground handling servicing, compliance with procedures, standards, SOP's emergency procedures and other activities**

Here follows an impact-relations map (IRM), with factors neatly organized in natural groups.

[14] Developed in 1971by the Science and Human Affairs Program at research center of The Battelle Memorial Institute in Geneva.

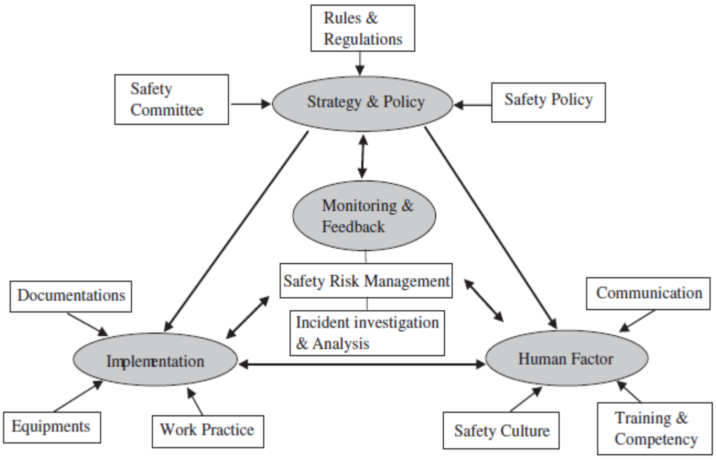

Here follows an Impact direction map, where the R+C on the horizontal axis is a measure of the total effect on the system for each factor, and R-C on the vertical axis is a measure of the net effect of the factor, where positive is interpreted as mainly a cause and negative as mainly an effect. The factors with large positive R+C on the high positive side is then taken to be the largest generator of effects, in as much as they affect other factors more than the other way around.

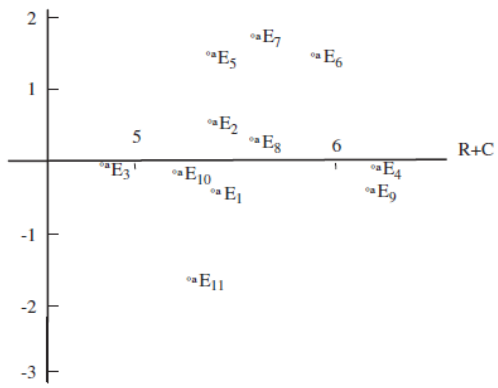

We see that the important factors turn out to be in this order E_6 (Rules and regulations), E_7 (Safety committee) and E_5 (Safety policy), which constitutes the natural group of Strategy & Policy on the top of the IRM map. These factors should then be the main area for improvement, in order to have a well-functioning safety management system (SMS).

Expert judgments in large projects: RPP

Large project involve typically many assessments of measurable quantities and events and their chances. With little data, we quite often have to rely on expert judgments. This is particularly so in the early project phase. It is a challenge to combine different judgments of the same quantity in a logical consistent manner. A variety of approaches have been suggested in the literature, but only few have obtained wide use in practice, most those already presented above. In recent years, one has realized that many of the quantities that are involved in large projects are correlated, so that judgments based on the assumption of independence will typically underestimate the risk. However, efforts to include opinions about dependence in the expert opinion have revealed several difficulties. It may be (i) more demanding for each expert to state an opinion, (ii) more demanding to summarize differing opinions and (iii) more demanding to plan and facilitate an expert opinion process. In some cases it may also complicate matters for the decision maker, but although it may be comforting to know that dependence have been an issue in the process. An interesting recent alternative is developed by a consortium of European universities financed by the EU commission. It is named Risk Planning Process (RPP) and described in more detail in Gasparini et.al (2004). Some characteristics of this approach: It allows updating during the various phases of the project. Every effort is taken so that the data demanded from the experts are something they can relate to in practice. Concerning correlation between events, only one of two is asked from the expert: Either high degree of correlation or not. With more than one expert one will hopefully get a balanced picture of the degree of covariation. Special attention is given to the assessment of rare events. The theory behind this approach is fairly advanced, but this may be hidden in user-friendly software.

The well-calibrated expert

Two questions of importance to risk management are:

1. How to obtain a fair opinion from expert?
2. How to combine opinions from several experts?

For an expert opinion on an event with several possible outcomes, we may want to have a statement on the probabilities of each outcome, summing to one. If the expert makes many such statements over time, we may in principle record the performance by comparing the probability statements with the actual outcomes. If there is a good match we would say that the expert is "well calibrated". The question is then: How to get a well-calibrated expert? We will briefly deal with this below. There are a number of possible biases and pitfalls in performing expert judgments in practice. Among them are:

- narrow scope
  – limited involvement
- mindset bias
  - adoption of own hidden assumptions
- motivation bias
  - experts own interests (often unconscious)
- cognitive bias
  - overoptimism,
  - anchoring: linking to less relevant ideas
  - accessibility: over- (under-) estimates the easily (hardly) accessible

There is an extensive literature on how to avoid them. It is difficult to be protected against all systematic biases, but knowledge about the traps may be useful. Quite often traps may be identified by a simple pilot study outside the expert panel, with no need to include the same expertise level as the panel. Some common probability biases are discussed in the next section.

A special problem is the assessment of rare events, and very few have the ability to relate to events with less than 1% chance to happen. Just because such events historically have been rare within the field in question, experts have had little chances to be good calibrated (and thus becoming experts). A possibility is to select some rare events from other fields or daily life, where the probabilities of outcomes are objectively known, and may order them. One may then ask the potential expert to order the outcomes according to the imagined likelihood to happen. This requires that the potential expert can relate to the events in the selected context. At the end one can choose to reveal the actual probabilities, so that the expert will have the opportunity to learn how to express the judgment numerically, even if the numbers are small.

Assume that the expert is asked to give his/her opinion on the outcome of a numerical variable X. This may be given in terms of a probability distribution F(x). However, in practice this may be to ask too much. Instead, we may be satisfied with some to ask for some quantile values, for instance the following: 0% (minimum), 5%, 25% (lower quartile), 50% (median), 75% (upper quartile), 95% and 100% (maximum). A p% quantile equal to x_p means p% sure that the value x_p is not superseded. In practice the following five may be sufficient 0%, 5%, 50%, 95%, 100%. The reason to omit the quartiles rather than the 5%, 95% quantiles may be that the extreme events are of more importance to assess. This divides the real line in four intervals having (subjective) probabilities resp. 5%, 45%, 45% and 5%. For calibration purposes, such a judgment is repeated n times for different unknown quantities, later to be known or kept secret. In each case we observe which of the four intervals given by the expert we find the actual value. Over the n repeats we have observed a fraction in each interval. If the expert is well calibrated, the fraction of hits should be in line with the four probabilities. A measure of how well the expert hits is based on the so-called relative information for the observed distribution with respect to the theoretical (5%, 45%, 45%, 5%) distribution. This measure is non-negative with maximal value of 1 if the expert hits perfect over the n repeats. In addition to be well calibrated in this sense, the expert may be measured on ability to provide informative statements.

An expert able to provide quantiles of his/her distribution, so that the range of possible outcomes can be restricted, is likely to be preferred. A measure of the expert's ability to provide information in this sense is also defined. A unified measure of expert "usefulness" is then given by the product of measure for "calibratedness" and "informativeness", and the expert's weight among several experts is determined based on this.

In many areas the decisions are made based on models with unknown parameter, which may not have a direct and intuitive interpretation. It may pose some difficulties to ask the expert to give their opinion on these parameters, even if they are used to think in terms of models, since the one in question may not be their favorite one. Most experts in field are likely to relate to observable quantities, and should preferable be given the opportunity to state their opinion in such terms, even if it may cause some trouble to translate this back to the parameters of the model. This may for instance be the case for issues related to environmental studies, where competing models are advocated, but may want to take advantage of a common expert base in the of scarce hard data.

Combination of expert judgments of probabilities

We end this section with a discussion of the combination expert judgment of probabilities, and give some examples to show that it is non-trivial to unify judgments by algebra only.

Given a prospect with m possible outcomes, and n separate expert judgments of the outcome probabilities. The objective now is to reduce the n ”competing” probability distributions to a single one to be used as decision support. A possibility is for each outcome to settle on an average of the n probabilities given to the outcome. We have several possibilities, arithmetic or geometric mean, possibly combined with principles that guarantees some logical consistency, at least so that all probabilities are between zero and one and sum to one.

Example**:** Combined judgment 1
Two experts judge both that the probability that a system works A is 0.2. This gives a combined probability 0.2, whatever mean is used. Suppose further that both agree that there are two mutually exclusive possible causes for the system not working B and C. Expert no 1 believes that the probabilities for B and C are 0.5 and 0.3 respectively, while expert no 2 believes the opposite 0.3 and 0.5, so that the sum in both cases is 0.8. If we first combine the two probabilities for B and the two probabilities C by taking geometric means, we get in both cases $\sqrt{0.3 \cdot 0.5} = 0.387$. We then have three events A, B and C which are mutually exclusive, with ”probabilities ” 0.2, 0,387 and 0.387 summing to 0.974 and not to 1, as it should. The logical way to modify this is to divide each of the numbers by the sum, to give 0.206, 0.397, 0.397. Now we have a probability for A, which is different from the 0.2 they agreed one, not much different, but more so if they disagreed more on B and C. This may give unintended consequences in more complicated situations. On the other hand, if we had used arithmetic mean instead of geometric mean, this problem does not arise.

The judgment of independence is important in many situations. It is desirable that when two or more experts judge two (or more) events to be independent, then their combined probabilities should also reflect this. This is not possible, using arithmetic means, but geometric means is precisely the method that secures this.

Example: Combined judgment 2
The risk of fire in a gas pipeline is studied. For fire to occur there must be a leakage (A) while at the same time there is an ignition (B). Two experts agrees that the two events are independent of each other, so that the probability of fire is the product of the probabilities for A and for B. However, they disagree on the probabilities. Expert no. 1 claims that they are resp. 2×10^{-4} and 4×10^{-4}, while expert no.2 claims they are resp. 8×10^{-4} and 1×10^{-4}. Here both agree that the probability of fire is 8×10^{-8}. If we combine their probabilities for each of the two events separately using arithmetic means, these become 5×10^{-4} and 2.5×10^{-4} which multiplied together gives 12.5×10^{-8}, which is more than 50% more than 8×10^{-8}, which was their individual consistent judgment, based on the common agreement of independence. If we instead use the geometric mean, the combined judgments of the separate probabilities 4×10^{-4} and 2×10^{-4}, which multiplied together gives 8×10^{-8}, so that their agreement on the fire risk is kept.

We see that one attractive feature is obtained by the arithmetic mean, while another is just obtained by the geometric mean. You cannot obtain both at the same time. Independence is by many regarded as an important feature that should prevail in the combined judgment when the experts agree on that is how it is. This suggests that we should use the geometric mean and rescale as in the example. On the other hand, it is claimed that this in some sense will be contrary to learning from experience, and therefore stick to the arithmetic mean, unless special circumstances points to a different choice. The remaining issue is then the weighing. The simplest is to use the ordinary equally weighed mean. However, the decision maker may want to put more weight on the expertise that has shown to be "more in line with reality". Experience on accuracy of the individual expert or expert field, may then in some cases justify another weighing. A scheme for performance based weighing, in cases where the experts have to express the probable values of a measurable quantity, is given by Cooke (1991), see also Bedford & Cooke (2001).

Bayesian principles for the combination of expert judgments may be applicable in many contexts, in the judgment of measurable quantities, as well as in the judgment of probabilities. If the judgments are part of a larger context as one of many, for example as part of a decision analysis the challenge is bigger.

2.8 Environmental risk strategies

In this section, we present some ideas that may be helpful when addressing some of the major human concerns, with consequences possibly far I the future. First, a method to derive solutions based on four different conceptions of the relationship between nature and society. Then, so-called robust decision making (RDM), which is a way to generate good solutions in complex situations. That is, we may have a wide set of opportunities, and the optimal choice will depend on some future state of nature, which we do not know, and is hard to model.

Conceptions of Nature and Society

Change typically affects both nature and society, and the appropriate strategy will depend how we interpret this relationship for the current issue. Consider the following conceptions:

1. Nature is in a fragile state of balance, which may easily and suddenly turn into another state where the some basic processes as we know them cease to function, and irreversibly so.
2. Nature aims at balance and even major disturbances are adsorbed to regain its former balance as we know it.
3. Nature may be in balance in several possible states, and in another state where something is lost, provides at the same time a new opportunity.
4. Nature has tolerance against disturbances, and they must last long to have major impact.

It is hardly possible to know for sure which conception is the true one, and one may wonder whether science ever will be able to answer this question definitely. For society we may also have different conceptions, perhaps similar to the ones for Nature. In our context we have to consider the effects on societies of nature changes and interventions. We may have better grasp on what is reasonable for a given society, but will not know for sure.
It may be useful to think in terms of dichotomies:

A. Nature may be robust or fragile to interventions, and
B. Society may be robust or fragile to interventions in nature with social consequences.

This gives rise to four combinations, for which the precautionary principle may point to different strategies, as outlined in the following table:

		Nature	
		Fragile	Robust
Society	Fragile	Clean technology strategy	Technologist strategy
	Robust	Deep ecology strategy	Policy control strategy

Note that all four conceptions may accommodate the view that we have a major environmental problem, and that we have to do something about it.

Let us look into how the precautionary principle may lead to the different strategies depending on the conception we take as our platform.

Deep ecology (or similar) strategy
With this platform, we will not take the risks by interventions to a presumed fragile nature, but are willing to take risk by intervention in human behavior, presuming that society is robust and will take it without malfunctioning, even if the social structure changes.

Clean technology strategy
With this platform, we have to act as if both nature and society are fragile. This gives little room for radical strategies, as typically something that may give radical improvement on a parameter related to nature, may have large effects on society. Rather one would choose strategies that are to some extent sensible whatever the later state of the world. One accepts that the burden on nature is unacceptable in the long run, and may then choose to adjust a lesser fragile parameter, technology. The idea is to develop new and more clean technology that fulfills the same socio-economic function as the old one. The point is to keep the burdens on nature and society acceptably low at the same time. This will typically correspond to a so-called regret-type strategy.

Political control strategy
With this platform, both nature and society are robust. Implicitly this means that both have wide ability for adjustment, and that we are not necessarily in a hurry. Long run and significant changes may therefore be achieved by a step-by-step political process, without excessive costs and without putting extra burdens on society (even if our judgment is that more radical actions would not disrupt society). The strategy then is political control by regulations, punishment and incentives. An important issue in control is the use of tolerance limits, which may work contrary to the precautionary principle.

Technologist strategy
With this platform, nature is robust and society fragile, so we cannot take the risk to implement drastic measures. Since nature is assumed robust, we probably have some time. The reasonable strategy would be to repair the current damage, as far as it is economically and technically possible, and then consistently push technology in the right direction, and at all times use the best available technology (BAT), and think in terms of ALARP ("As Low as Reasonably Practicable"). Note that although this has some resemblance with clean technology, the time perspective is different.

Discussion:
Suggest possible strategies/policies for each of the conceptions of Nature and Society for each of the following threats, and discuss their pros and cons.

a. Fish farm escapes – a threat to the wild salmon population
b. Salt nutrients to the North Sea – a threat to sea life
c. Global warming – a threat for everybody
d. Barents Sea Drilling – a threat to the Arctic?

Robust decision making (RDM)

This is a strategy for providing decision support in case of deep uncertainty and strong disagreement about assumptions and values, as is often the case for environmental issues.[15] RDM has three main characteristics:

1. Multiple views of the future
   — RDM rejects the view that a single joint probability distribution is the best description of a deeply uncertain future. Rather RDM uses ranges or sets of plausible probability distributions to describe deep uncertainty
2. Robustness rather than optimality
   — RDM judges alternative strategies by a robustness criterion rather than an optimality criterion. Instead of trying to determine the highest ranked option in some sense, for example by the expected utility criterion, RDM will be based on some chosen satisficing criterion, and then describe tradeoffs judged by chosen robustness criteria. By RDM one is willing to trade a small amount of optimum performance, in order to obtain good performance over a wide range of plausible scenarios, and keeping options open. An RDM strategy will typically be less sensitivity to broken assumptions than traditional (expected utility) decision strategies.
3. Uncertainty judged in the context of a particular decision
   — RDM first identifies the combinations of uncertainties that are most important to the choice among alternative options, and then describes the set of beliefs about the future state of the world that are consistent with the choice of one option over another. Traditional decision analysis, on the other hand, first characterizes uncertainty about the future without reference to the decision options, and then uses this as basis for ranking the alternatives options. The RDM approach may allow the stakeholders to understand the key assumptions underlying alternative options before committing themselves to believing those assumptions.

RDM will typically be an iterative process based on a "vulnerability-and-response-option" framework rather than the common "predict-then-act" decision framework. Often, these strategies are adaptive and designed to evolve over time in response to new information. Robust decision methods seem most effective when:

— the uncertainty is deep and not well-characterized,
— there are many decision options,
— the decision situation is complex

RDM is not a recipe, but rather a set of methods that can be combined in varying ways. Central to RDM are

- Exploratory modeling
- Scenario discovery

The exploratory modeling typically makes use of simulation models to trace the potential consequences over many plausible scenarios, not for prediction, but rather for relating

[15] Largely developed within the RAND corporation.

assumptions to their implied consequences. Repeated runs may be performed under different assumptions, covering the range of parameter uncertainties, often within an experimental design framework. This establishes a large database of result cases, which can be analyzed to identify vulnerabilities of proposed strategies and the tradeoffs among them.

A workbench for exploratory modeling in support of robust decision making under deep uncertainty, may be found in http://simulation.tbm.tudelft.nl/ema-workbench/contents.html.

Scenario discovery is a process to aid the identification of vulnerabilities linked to the proposed strategies. The process requires some performance metric, such as the total cost of a strategy or its deviation from optimality (regret). This is then used to distinguish those cases in the result database where the strategy is judged successful from those where it is judged unsuccessful. Statistical or data-mining algorithms may then be used to describe the regions of the input parameters space that best describe when a strategy is unsuccessful. Thus the algorithm is tuned to take into account both the predictability and interpretability by decision-makers.

2.9 Probability and risk literacy

Risk fallacies

People may find it hard to evaluate probabilities, and investigations have shown that evaluations quite often are biased and inconsistent. Research into this and how people go about evaluating probabilities have been pioneered by Kahneman and Tversky.

In case of no scientific (empirical) knowledge or any method to reveal probabilities scientifically, we have to rely on experience and introspection. However, experience may fool you for several reasons:

- Experience is observing events under circumstances far from scientific
  - It is selective and open to individual interpretation.
  - What we conclude and choose to remember may further depart us from reality.
- Without feedback on past actions related to the current, we cannot trust experience.

Most people seem to be overconfident. This is confirmed in studies where the subjects are asked to make a large number of predictions, and at the same time state their prior beliefs on the chances of the predictions come true. It turns out that errors of prediction occur more often than their stated prior beliefs, and quite often more than just marginally worse. A noncalibrated person may typically err about 65% of the time when claiming to be 90% sure. Of course some appear very optimistic in order to get going something, say a project they have stakes in, but overconfidence seems to be prevalent, even with no stakes.

The human mind is not a computer and typically has to rely on heuristics, that is, mental short-cuts that make it possible to cope with the realities of this world. When dealing with risks heuristics quite often lead to biases. Here are some typical fallacies:

<u>Randomness vs. Representativeness</u>: People tend to judge clustered patterns as non-random. For instance when comparing the results HHHTTT and HTHTTH for six flips of a coin, although both have the same probability 1/64. It is just that there are many more mixed patterns. Similarly may a cluster of accidents be taken as increased objective risk which requires action, which may be false, since random instants will typically not appear as a regularly spaced pattern. Another example is the citizens of London during World War II who felt that certain areas were targeted, when in fact the pattern was more like a random one.

On the other hand, a long run of no adverse event may often be taken as a system improvement. However, occasionally long runs are consistent with pure randomness and no system change to the better. This may lead to a more relaxed attitude, which may rather increase the risk. This may also happen in the case when several near misses are observed, since the fact that the miss did not happen may falsely be taken as the system is robust against misses. However, after a major miss many are eager to state that they saw it coming all along. These issues are dealt with in Statistical Process Control (SPC).

<u>The conjunction fallacy</u>: People often tend to violate the necessity that if A implies B the probability of A has to be at most B. An example is air travelers who are willing to pay more for an insurance policy that covers terrorism, than a policy that covers any cause of death during the

flight. This may be explained by a tendency to pay more attention to the most specific cause and attribute highest probability to that, regardless of the logical inconsistency.
The small number fallacy: Suppose you have a method for performing a difficult task, said to be successful about 70% of the time. You try it out 12 times with 8 successes and 4 failures. Then someone suggests a new method, and you try it out 4 times with 0 failures. With just this information it is likely that many of you may claim the new method is regarded as the superior one, despite the fact that the probability of the observed result may be of the same magnitude for the two methods.

Ignoring variance: Quite often an issue is discussed based on an average quantity computed from a sample from a population where the population mean is of prime interest. In this context we have at least two traps:

- Ignoring that the variation in the population may be relevant in itself.
- Ignoring the fact that the computed average is just an estimate of the population mean, This is worst in the case of small samples, which frequently is the case in the media.

Insensitivity to prior knowledge: The classical example is a screening test for a serious disease, which does not give 100% correct answer, in that we may have both false positive and false negative. If the prior probability (without any test) of having the disease is low, we may have the situation that the majority of those found positive by the test (i.e. indicative of the disease) are in fact negative.
Exercise: Take $P(+|Sick)=0.90$, $P(+|Healthy)=0.05$ and prior probability $P(Sick)=0.01$, and compute $P(Sick|+)$ by Bayes law.

Many misinterpretations are related to conditional probability statements, among them

Fallacy of the transposed conditional: $P(A|B)$ is confused with $P(B|A)$.
Example: "If the accused is innocent there is a 10% chance of having the blood type found at the crime scene. Consequently it is 90% chance that he is guilty".
With B=matching evidence and A=Accused innocent, the statement $P(B|A)$ is falsely taken as $P(A|B)=0.10$ from which $P(A^c|B)=1- P(A|B)=0.90$ is obtained. In the current context the fallacy is named " the prosecutors fallacy".

Risk communication

Statements about hazards are often misunderstood by the public, by the media, and even by professionals in the field in question. The way risk is communicated is therefore of importance. In some cases the communicators deliberately use the phrasing to misguide the public.

We will consider three forms of risk communication that may go wrong:

- Single event probabilities and context ambiguity,
- Risk reduction: Relative risk vs. Absolute risk
- The benefits of testing: Conditional probabilities.

Reporting single event probabilities:

A single event probability statement is one of form:

"The probability that the event A will happen is p%"

This statement can be confusing without a proper reference context, and may mistakenly be taken as more than just an uncertainty statement, which may have no frequency basis at all.

Example: "The chance of rain tomorrow is 30 percent" is a probability statement about a one-time event. It will either rain or not rain tomorrow, but the statement can never be proven wrong whatever happens. The statement is ambiguous, and some people may interpret it different from what the weather forecaster have in mind. Some may think the statement means "it will rain 30 percent of the time", others that "it will rain in 30 percent of the area the forecast is given for", and finally, some believes "it will rain on 30 percent of the days that are like tomorrow". In case of ambiguity, people tend to choose their own reference context, but here the last interpretation is the one closest to what weather forecasters have in mind. By contrast, the statement "it will rain on 10 days in August" can be proven true or false, since it is not a one-time statement, but a frequency statement.

Reporting the benefits of treatment: Relative risk vs Absolute risk:

Consider the following newspaper headline:

"By taking drug X the patients with Y will reduce their of dying within 5 years by 22%"

This may look impressive, but what does 22% really mean? Studies have shown that many people mistakenly interpret this as out of 1 000 Y-patients, 220 of the deaths within 5 years can be prevented by taking X. Suppose the results of the study were as follows[16]

	Died	Not died	Total
Treated	32	968	1000
Not treated	41	959	1000

We see that 32 died within 5 years among 1 000 Y-patients who took X, while 41 died among the 1000 Y-patients who got a placebo. This may be communicated in several different ways which may trigger entirely different emotions. The headline statement is the *relative risk reduction*, obtained by (41-32)/41=0.22. On the other hand we could report the *absolute risk reduction*, which is the proportion of patents who die without treatment minus those who die with treatment (41-32)/1000 = 0.09. However, a headline saying 0.9% risk reduction of certainly does not as impressive as the 22% reduction of the headline above. In reporting to the public absolute risks should be preferred to relative risks, as it is more likely to be understood correctly. However, a good alternative is the *number needed to treat (NNT), defined by* the number of people who must participate in the treatment to save one life. In our case NNT=111, which is obtained directly from the absolute risk reduction as 9/1000 = 1/111. A statement like "Out of 111 treated patients, 1 had the benefit of the treatment, whereas the other 110 did not", is most likely understood.

[16] Numbers from a Scottish study on coronary heart disease and a cholesterol-lowering drug named Pravastatin.

Reporting the benefits of testing:

Consider a test that is supposed to give an indication of whether an unwanted state exists or not. A good example is a screening test to detect a disease.

"*If a woman has breast cancer, the probability that she will test positive on a screening mammogram is 90%".* Quite often people will confuse this with the statement: "*If a woman tests positive on a screening mammogram, the probability that she has breast cancer is 90%".*

This is another example of the transposed conditional fallacy mentioned above: The conditional probability that an event A occurs given event B is confused with the conditional probability that an event B occurs given event A. One can reduce this confusion by giving the statement in terms of natural frequencies instead of probability percentages.

Spurious ranking

Ranking performance is done in many areas, among others by comparing the failure rate on tests for schools and medical treatment among hospitals. Often results are brought to the media with the intention to help the public to choose among alternatives or question why a specific school or hospital is doing worse than others. Quite often unjustified statements are made. First we have the issue that the comparison may not be fair. Some schools have more pupils that need special attention and some hospitals may get the patients with poorer prognosis in the first place. Then, even if such differences do not exist, the ranking may be spurious. In a pool of schools or hospitals someone has to be on top and some on bottom due to pure randomness, and those who report the results quite often do not take this into account at all. For this reason it is advocated that results on studies like this should not be communicated as rankings. A good alternative is a so called funnel plot.

Example

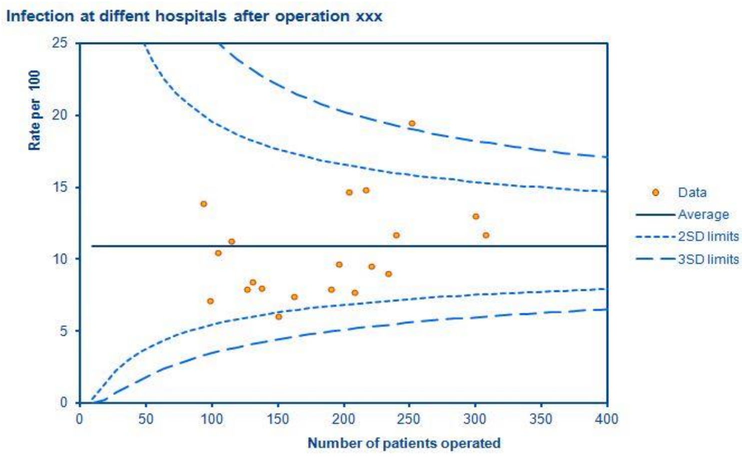

”The 100 year flood”

In their commentaries to an extreme flood events the media often refer to ”the 100-year flood”. What is a meaningful interpretation of this notion? First, it should refer to a well defined event at a specific location, e.g. a specified tidal height at a location protected by dykes. A “100-year event” sounds like an event we expect to happen about every 100 year (on average). However, many misinterpret this, and think in terms of a fixed time horizon of 100 years. This may lead to statements like:

(i) “No 100-year flood has happened so far in 90 years, so we are sure to get one soon!”

(ii) “We got the 100 year-flood last year, then we are safe for many years to come!”

(iii) “We got two 100 year-floods in this decade, so the risk of floods must have increased!”

These are all statements that do not hold any critical scrutiny, and you hopefully agree. Nevertheless, it turns out that politicians, planners and supervisory authorities sometimes plan or act as if there was a fixed 100 year horizon to judge an safety investment. In the phrase “a 100-year event is expected to happen about every 100 year on average”, the words “expected” and “on average” are crucial, pointing to a statistical and probabilistic interpretation. A possible interpretation of “a 100-year event” is an event having a probability of 1/100 of occurring in any given year. Alternatively, we interpret 100 as the expected number of years until it occurs, starting from any given year, regardless of the history. In fact, under reasonable assumptions, the second interpretation follows from the first. These interpretations are both reasonable and useful, and let us look into this more closely:

Assume that the extreme event A is well-defined, and that we observe at the end of each year whether it has occurred or not within that year. We assume that the probability of A happening in any given year is p, and the same from year to year. If p=1/m where m=100 we may talk about a 100-year event. Starting from year 0, we may think of the (random) waiting time N (in number of years) until the extreme event A happens. Assuming that the outcomes in subsequent years are independent, we have that the probability of waiting time equal to n is given by

$$P(N=n)=(1-p)^{n-1}p \qquad \text{for } n=1,2,\ldots\ldots$$

This is the so-called Geometric distribution. For this distribution we have that the expectation of N is E(N)=1/p, which in the case of p=1/100 gives E(N)=100.
We see that the range of possible values of N is unlimited. Moreover the distribution is without memory, in the sense that for any k>0

$$P(N=k+n|N>k)=P(N=n)$$

Which says that whatever how long we have waited without seeing the extreme event A happen, the remaining waiting time probabilities are exactly the same as they were initially. This means that the expected waiting from year k on is still 1/p, i.e. 100 years for a 100-year event. For a given time horizon h we may, with the assumptions taken, experience the extreme event once, more than once or never. In fact the number of extreme events X in h years is distributed binomial(h, p)

$$P(X = x) = \binom{h}{x} p^x (1-p)^{h-x}$$

In the case of h=100 and p=1/100 we have

x	0	1	>1
P(X=x)	0.366	0.370	0.264

Another question to be asked is: How to define a "100-year event"?
This may be done by observing sufficiently many yearly maxima, and establishing the distribution, and the compute its 99% upper quantile. This may be purely empirical, or fitting a suitable parametric distribution, possibly derived from available theory in the area of application, or general extreme value theory.

Remark. The notion of 100-year flood was introduced in the 19060's by the US Geological Survey. Due to the confusion among laymen, the now advocate its replacement by the notion Annual Exceedance Probability (AEP).

Exercises

Risk calculations may be flawed for various reasons. In each of the statements below, point out the inherent danger of having understated the risk.
(Answer each in maximum 3 lines).

(a) "The hazardous event B is triggered only if both events B_1 and B_2 happen. Since each of them happens with probability 0.01, B will happen with probability 0.01· 0.01=0.0001".
(b) "The hazardous event B is triggered if one of the events B_1 or B_2 happens. Since each of them happens with probability 0.01, B will happen with probability 0.01+ 0.01=0.02".
(c) "For our product xyz, the sales (X) for the next month is judged having expectation 1000. We sell at market prices (Y) expected to be 5. Then the expected income from xyz is 5·1000=5000. Even if sales turn low due to adverse weather conditions, I don't think we are in trouble".
(d) "We have diversified by investing evenly in four stocks, each with expected return 5% and standard deviation 10%. Thus our portfolio is expected a return of 5% with standard deviation cut to 5%".
(e) "Mrs. A: I have a well diversified portfolio of stocks giving 10% expected return with standard deviation 5%, so that the chance of getting negative return is just about 2.5%". "Mr. B replies: I have a single investment giving the same expected return and same standard deviation. So I achieve the same. Why bother with all those different stocks!"

3 Special analytic topics

The third part of these notes contains some useful analytic topics, which may be read independently of each other.

3.1 Classes of useful distributions

Some useful classes of distributions are:

Discrete distributions	Continuous distributions
Binary	Normal
Binomial	Lognormal
Poisson	Exponential
Geometric	Gamma
Negative binomial	Weibull
Hypergeometric	Beta
Uniform	Uniform (Rectangular)

Brief description of some of them:

Discrete distributions:

Binary distribution (also named Bernoulli-/indicator- distribution): Supports the values 1 or 0 with probabilities p and 1-p respectively, e.g. a indicator for events ”failure”=1 and , ”no failure”=0. Notation: Bernoulli(p)

Binomial distribution**:** represents the number of 1's for n independent repeats of the same binary variable, e.g. the number of failures among n trials. Notation: Binomial(n.p)

Poisson distribution**:** Supports a non-negative integer variable, with no upper limit. May fit to situations where events occur in (continuous) time independent of each other, e.g. the number of accidents in a given time interval. Notation: Poisson(λ), where λ is the expectation.

Geometric distribution**:** Supports a positive integer variable, with no upper limit. An example is the time until the first failure in a sequence of repeats with failure probability p. Notation: Geometric(p)

Negative binomial distribution**:** Similar to the geometric for the time until the k'th failure.

Continuous distributions:

We have mainly three types of continuous distributions: (1) unbounded two-sided support, (2) unbounded one-sided (typically positive) support and (3) bounded support.

Here are examples of one distribution of each type:

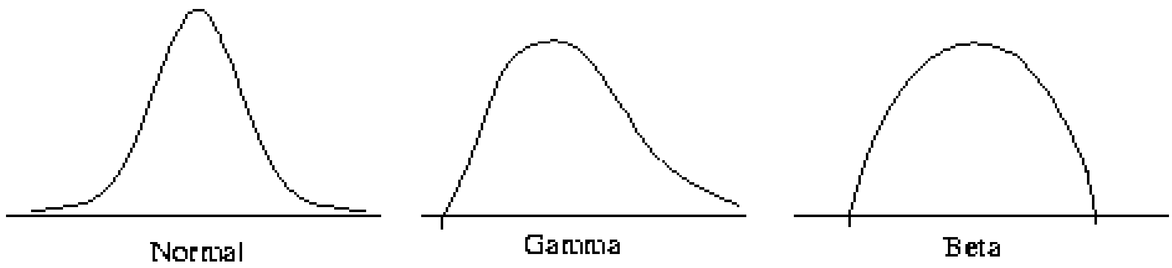

Normal distributions: Symmetric about an expected value μ with deviation risk measured by the standard deviation σ, and so that deviations from the expectation more than +/- $k \cdot \sigma$ for k=1,2,3 are 32%, 5% and 0.2% respectively. Notation: Normal(μ, σ).

In many risk problems extreme deviations are more frequent than the normal distribution can account for, and a symmetric distribution with heavier tails is more appropriate. Distributions of this kind are the t-distribution, the logistic distribution and the most extreme of them all, the Cauchy distribution.

Lognormal distributions: Positive support, and a unbounded right tail, and may also be specified by its expectation and standard deviation. Often used as model for losses.

Gamma distributions: Also positive support and a unbounded right tail, but are more varied with respect to shape than the lognormal distributions. Often used as model for loss and spent time. Notation: Gamma(α,λ) where α is a shape parameter and λ is a scale parameter. Gamma(1,λ) is the exponential distribution.

Weibull distributions: Also positive support and a unbounded right tail. Often used as model for life times. Notation: Weibull(α, λ) where α is a shape parameter an λ is a scale parameter. Weibull(1, λ) is the exponential distribution α=1 corresponds to the exponential distribution (constant failure rate), while α>1 and α<1 correspond to increasing and decreasing failure rate, respectively.

Beta distributions: Support the interval [0,1]. Often used as model for fractions, and may provide symmetric as well as skew cases with both small and large risks for deviations from expected value. Notation: Beta(r,s), where r and s are the shape parameters. Here r=s=1 gives the Uniform[0,1] distribution.

Remarks

1. In addition to the distributions mentioned, we have the distributions occurring in statistics: Students t, Fishers F and the chisquare-distribution.
2. In some applications, the variable is restricted to the left by a number different from zero, e.g a minimum loss. We may the shift the distribution appropriately. Likewise, we may rescale the Beta distribution to any bounded interval [a, b].
3. For bounded variables the triangular distributions are also found useful in practice.. Notation: Triangular(a,b,c), where a is the minimum possible value , b the maximum possible value and c the most likely value.
4. In some areas, very heavy tailed distributions are needed. Some of this kind are: Pareto, Frechet, among others found useful for losses in finance.

Illustrations:

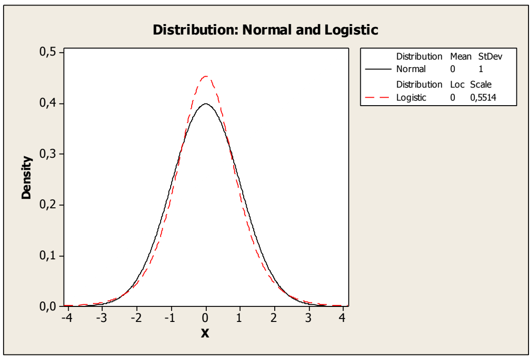

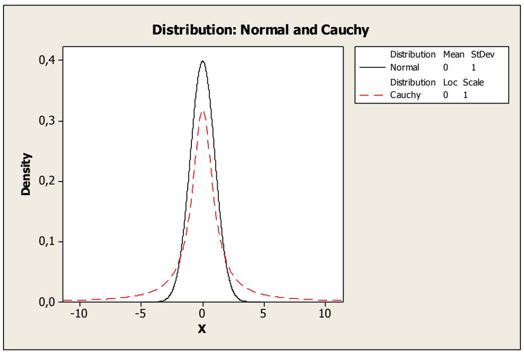

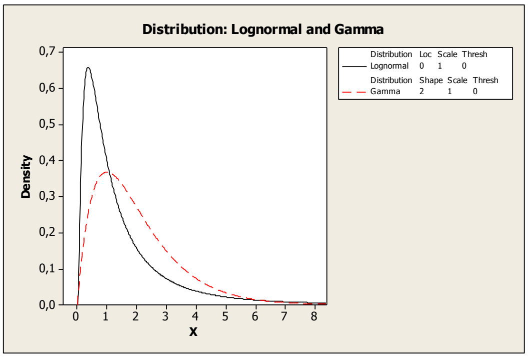

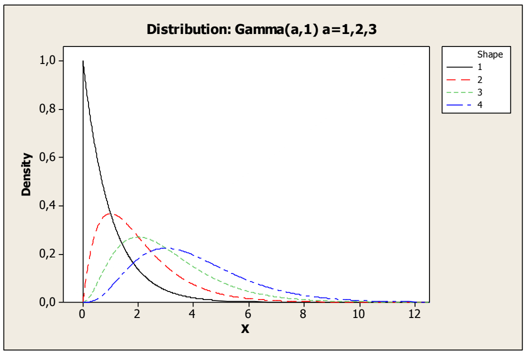

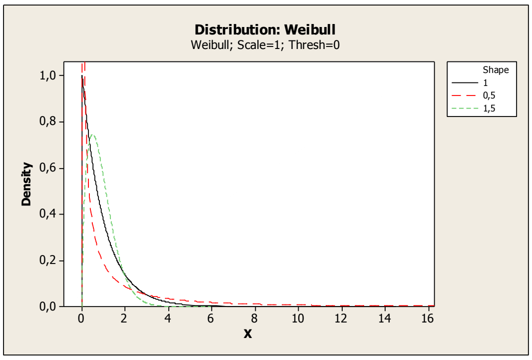

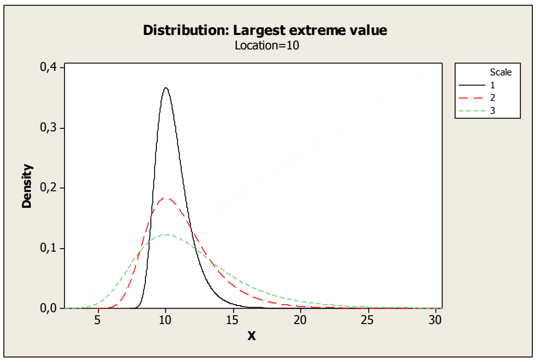

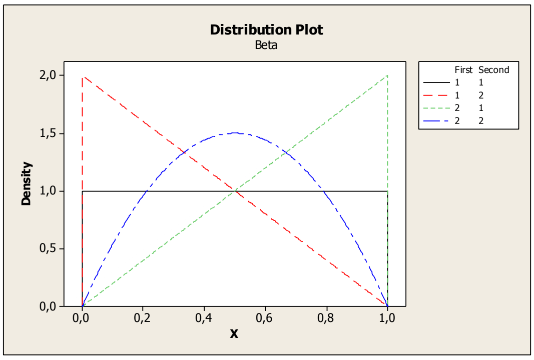

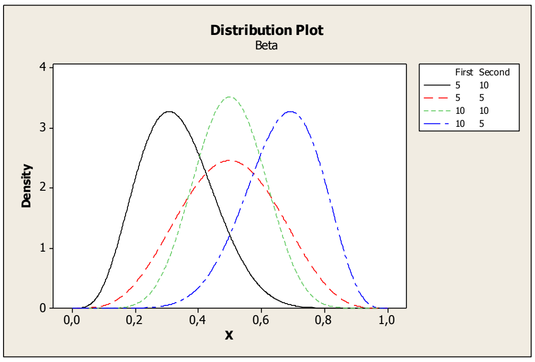

Example

A product has a component that needs frequent replacement, and the choice is between two brands, a new brand (1) and the old (2). The life lengths may vary between components, even within brands, and which one is the best is tested in an experiment where the life length of 50 units of the new and 40 units of old is observed. However, the experiment cannot go on forever, so at about (life) time 70 those who are still working are recorded with their current age (right censored data). We got the following:

	# observed lives	# censored lives	# total observed
Brand 1	37	13	50
Brand 2	34	6	40

The actual numbers are depicted in the following graph (censored units in black):

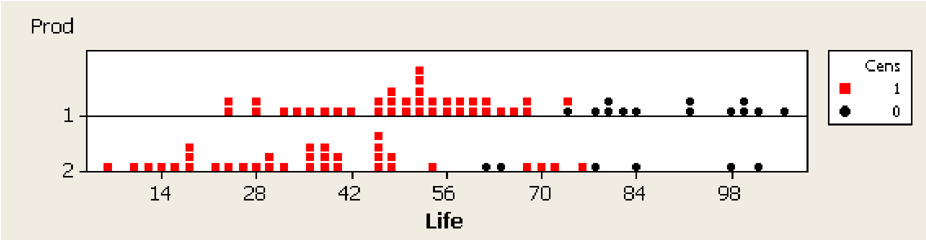

Here follows output from an analysis based on the Weibull life model, where the information given by the censored observations are taken into account:

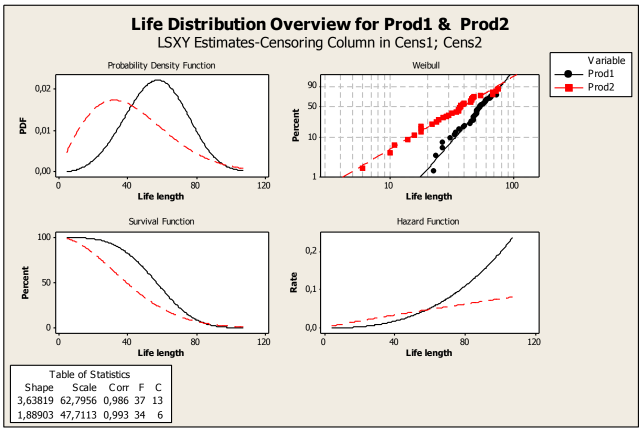

In the panels are: The estimated life densities, the survival function, the hazard rate functions, the mode fits, and the estimated Weibull distribution parameters. For details on the Weibull distribution, survival function and hazard rate function, see the section on process theory.

3.2 Sampling inspection

Sampling inspection based on statistical principles is useful in many contexts in industry and elsewhere. We may have sampling for

- Attributes (e.g. quality in ordinal categorical classes)
- Variables (e.g. quality measured on a continuous scale)

Sampling schemes fall in two broad categories

- Sampling from lots (i.e. finite population)
- Sampling from process (i.e. independent repeats)

We will limit the discussion here to attribute sampling with two classes, here named defective (d) and non-defective/intact (i).

Sampling for defective in process

Let

p = the probability of an item being defect
n = number of produced items (sample size)
X = number of defectives in the sample

If the outcomes of the produced items are independent repeats, the probability distribution of X is Binomial(n, p) i.e.

$$P(X = x) = \binom{n}{x} p^x (1-p)^{n-x} \quad x = 0,1,...,n$$

with expectation and variance given by

$$EX = n \cdot p \text{ and } \operatorname{var} X = n \cdot p \cdot (1-p)$$

If p is small and n is large:

X is approximately Poisson(np) : $P(X = x) \approx \frac{(np)^x}{x!} e^{-np}$

When np > 5 the approximation is typically correct to the second decimal, which is sufficient in many applications. Example: p<0.10 and n>50.

Sampling for defective in lots

Let

N = number of items in population (lot size)
M = number of defective items in population
n = number of items in sample
X = number of defectives in the sample

For a random sample the probability distribution of X is Hypergeometric (M,N,n) i.e.

$$P(X = x) = \frac{\binom{M}{x} \cdot \binom{N-M}{n-x}}{\binom{N}{n}} \qquad x = 0,1,\ldots,\max(M,n)$$

with expectation and variance given by

$$EX = n \cdot \frac{M}{N} \text{ and } \operatorname{var} X = \frac{N-n}{N-1} n \cdot \frac{M}{N} \cdot (1 - \frac{M}{N})$$

If N is large in comparison with n:

X is approximately Binomial(n, a=M/N)

If furthermore a is small and n large, but still small in comparison with N:

X is approximately Poisson (na)

Remark.
Note that sampling for defective in lots presuppose that defectives are produced to some extent. In modern manufacturing the fraction of defectives are so small that sampling for defectives are like search for "the needle in the haystack", and will not be cost-effective. Moreover, the production of defectives is a waste in itself, and should be avoided. A helpful tool to achieve this is statistical process control (SPC). With this we can detect changes in the production process, which may lead to defects at an early stage, before defects are produced. It is said that those who have to rely on sampling for defectives in lots will soon go out of business in the competitive environment of today.

Exercise

The management of a fish farm for salmons ("salmo salar") fears that the location is infected by the Salmonid alphavirus (SAV) leading to pancreas disease (PD) and premature death. Even if just a few salmons are infected, this may easily spread all over, with large economic consequences. . The management therefore plan to take a sample from the enclosures (named "merds") to investigate if the virus can be detected at all among the estimated 800 000 salmons in the merds. The question is how big a sample?

3.3 Statistical process control

Statistical process control (SPC) is a valuable tool for quality improvement in many industries and may be used in services as well. It is well described in the quality improvement and management literature. In fact, it is more than a tool. It plays a central role in some management theories, most notably Six-Sigma. It is somewhat peculiar that it does not play a larger role in the risk management literature. The basic elements are as follows:

1. A process view on all activities
2. Understand the variation in products and processes
3. Reduce variation as a key objective
4. Measure key characteristics of products and processes
5. Chart the measured characteristics in control charts
6. Take action according to basic understanding of the nature of variation

It is fruitful to imagine two types of variation:

- Common cause variation
- Special cause variation

Common cause variation is the variation inherent in the process, typically due to many not directly identifiable causes, while a special cause variation is the variation due to some specific cause, and not really inherent in the process. Special cause variation can in principle be removed from the process by identifying the special causes and remove them. This is something that can be delegated down in the organization, in production to operator level, given the required understanding and tools. Common cause variation cannot be reduced without changing the process itself, which is a management responsibility. To uncover opportunities for variation reduction requires a deeper understanding of the process.
In quality management terminology, the handling of special causes is taking corrective action, while common cause variation reduction is about quality improvement.

Production must relate to specifications and tolerances. There are two views on this

– Everything produced within the tolerances are equally good.
– Every deviation from nominally value represent a loss, more so the further you are off.

In many industries, the first view has been predominant, and maybe still is. With this frame of mind, there is no incitement to improve as long as you are within the tolerances. However, those who adopt the second frame of mind, will typically experience that their improvements will be known in the market, and then they are ahead.

In order to get improvements the enterprises have to focus on key characteristics, which is features or characteristics whose variability have the greatest impact on the performance for the customer. The chosen key characteristics then have to be measured over time, and then points are plotted in a control chart, in order to reveal the kind variation present.

There are two possibilities:

1. The process is “in statistical control”: The points exhibits a random behaviour over time around a central line
2. The process is “out of statistical control”: The points exhibit non-random behaviour of some kind: extreme single points, level shift or trends.

In the first case we interpret the variation as common cause variation. In the case that there are also special causes present, they are indistinguishable from the common causes, and wasted time to chase. In the second case we interpret the picture as there is special cause variation present, not inherent in the process, which may be identified and removed, and then get a process in statistical control.

Here is a control chart for a measured quality characteristic in statistical control.

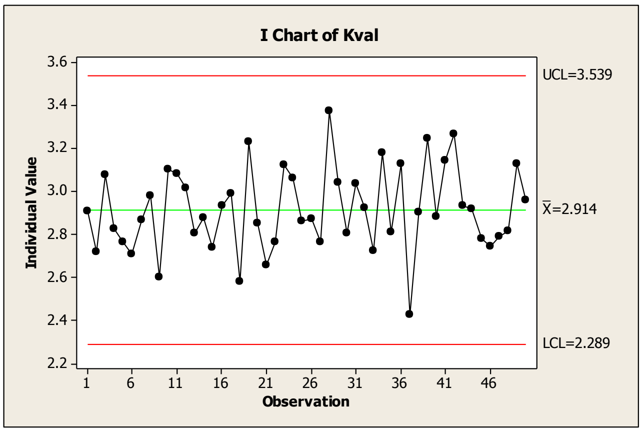

The points for 50 periods are plotted in the order observed, and the pattern looks random around the centreline. Common practice is to compute control lines at plus/minus three standard deviations from the centreline. The established lines may then by used for further monitoring the process. Here in a continuation for 20 more periods.

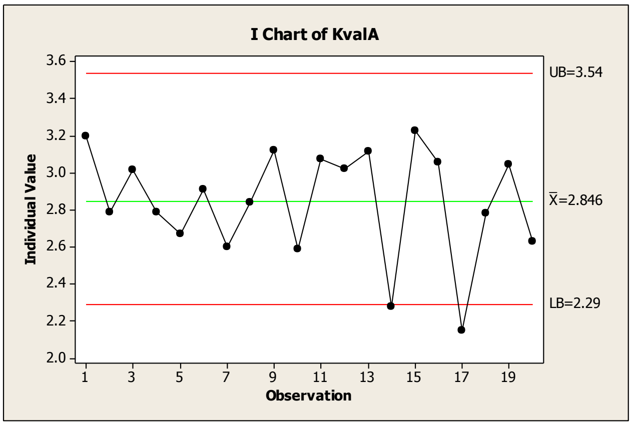

We see that the process is out of control at observation no. 14 and 17.
With a process out of control, anything can happen, and risk cannot be calculated. Moreover, such a production or service regime may require costly end inspection. With a process in control guarantees may be given, and in many cases, no end inspection is required. It will therefore be of prime importance to bring the processes measured by the key characteristics under statistical control, and monitor them by control charts. To bring a process in statistical control we may have to look upstream for causes, and perhaps start measuring some process characteristic there.

Control charts may also be used to judge whether a change, intended to improve the process, has really achieved variation reduction

There are a number of different types of control charts available, the main categories are attribute charts (e.g defect counts) and variable charts (measured characteristic, as above).
In practice one has to deal with a number of questions: How frequent should we observe? Should we measure in batches, e.g. 5 units and plot the average? We will not go into these issues.

Control charts may be useful in many areas for monitoring, among them

- Industry: Production - act on defects and process changes
- Health: Birth malformations, infection at hospital ward etc.
- Administration: ...... imagine yourself!

3.4 Active monitoring of accident data

Many companies and organizations observe and keep records of work accidents or some other unwanted event over time, and report the numbers monthly, quarterly or yearly. Here are two examples, one from a large corporation and one from the police records of a medium sized city.

Month	1	2	3	4	5	6
#Accidents	2	0	1	3	2	4

Year	2000	2001	2002	2003	2004	2005	2006
#Assaults	959	989	1052	1001	1120	1087	1105

In both examples the numbers show a tendency to increase over time. However, the question is whether this is just due to chance. If we add one to each accident count, the numbers do not look much different from those coming from throwing a fair dice. Taking action on pure randomness, believing that something special is going on is, at best, a waste of time. Randomness, or not, may be revealed by observing a longer period of time, but there may be no room for that. Pressure will rapidly build up to do something. In the case of accidents, find someone responsible. However, this may lead to blaming someone for something that is inherent variation "within the system". In the case of assaults, media attention, where often just this year is compared with the last or some favourable year in the past, leads typically to demands for more resources or new priorities.

Can we get some help from a probabilist or a statistician? The probabilist may say:
" If accidents occur randomly at the same rate, the probability distribution of the number of accidents in a given period of time is Poisson. If the expected number of accidents per month is 2, then the probabilities of a given number of accidents in a month are as follows:

#Accidents	0	1	2	3	4	>4
Probability	0.1353	0.2707	0.2727	0.1804	0.0902	0.0526

Thus the pattern of observations is not at all unlikely. There is even a 5% chance that you get more than four accidents, and this will surely show up sometimes, even if the system as such is the same".
The statistician may perhaps add: *"The average number of accidents over the six months is 2, but this is an estimate of the true expected number of accidents in a month. Typical error margins are plus/minus square root of 2 (knowing that the standard deviation of the Poisson distribution is the square root of the expectation), which is about 1.4. Thus the expectation may be anywhere in a wider region around 2, so that any claim that the 4 accidents last month are due to increased hazards is even more unjustified".*

You have taken a course in statistics in school, and have been taught about hypothesis testing, and you wonder whether this can be used. There are some problems: (i) It assumes a "true" expected rate and independent random variation, (ii) it typically compares expectations in one group of data with a hypothesized value, or with other groups of data. Here we have just instants, (iii) it may neglect the order of the data and may not pick up trends, and (iv) it does not handle the monitoring over time in an attractive manner. These objections may be remedied by more sophisticated modelling, but becomes unattractive for common use.

You may also have heard about statistical process control, and the use of control charts to point out deviant observations see (section 3.5). This keeps track of the order of observations and may also react to certain patterns in the data, like level shifts and increasing trends. The drawback is: (i) it assumes at the outset a stable system (a process in "statistical control") and the monitoring of deviant observations and patterns (ii) the data requirements are in many cases demanding. In practice, we often do not have a stable system, things change over time, from year to year, and we do not have enough data to establish the control limits, even if the system was stable.

Avoiding the reaction to deviant observations, which are just random, is a big issue in quality management, since it typically leads to more variation and turbulence in the system. This is justified in a context where you ideally have a system, stable over time, and running according to fixed procedures (until they are changed), e.g. in a production context and for some administrative processes. You may also have this in some contexts closely related to human hazards, say monitoring birth defects.

We probably have to realize that in many situations related to hazards there is no stable system and cannot be, and that some kind of reaction is required before we know the truth. What then to do? If we ask a risk analyst, he or she may be stuck in the traditional statistical theory, while others have dismissed traditional statistical theory as basis for active risk management (see Aven: Risk Analysis, 2003). Data of the kind above occur frequently, and there is room for a more pragmatic attitude. Otherwise the data will either be neglected, because no one can tell how to deal with them in a systematic manner (and particularly so the statisticians) or they will be misused for opportunistic purposes.

One possibility for monitoring with the aim to pick up trends in short series is described in Kvaløy & Aven (2004). It is slightly modified here, and is easily programmed in Excel:

Theory

A sequence of hazards for r consecutive periods is given, and the objective is to point out a worsening trend or individual hazards that are aberrant. Do the following:

1. Calculate the average m_j of the observed hazards up to and including period j
2. Take $e_j = (r-j)m_j$ as expected number of hazards for the remaining r-j periods
3. Use the Poisson distribution with expectation e_j to calculate the probability that the number of hazards for the r-j remaining periods is at least as observed
4. If this probability is small, say less than 5% then initiate warning or alarm.

Repeat 1-4 for all or some of j = r-1 , r-2 ,…,1

<u>Note:</u> If the counts are very high, we may use the normal distribution with expectation e_j and standard deviation the square root of e_j instead.

The first example above gives this calculation scheme for judgment at month six:

Month	1	2	3	4	5	6
#Accidents	2	0	1	3	2	4
Average till now	2.0	1.0	1.0	1.5	1.6	2.0
Expected ahead	10.0	4.0	3.0	3.0	1.6	-
Observed ahead	10	10	9	6	4	-
Probability (tail)	0.5421	0.0081	0.0038	0.0840	0.0788	-

We see that alarm is given at this month, due to some small tail probabilities looking ahead from month two and three. The 4 in the sixth month (observed ahead from the fifth) is not that surprising, judged from the average of the preceding. Neither is the combined 6 of two last months judged from the average of the preceding. However, the combined 9 of the last three months is dubious compared with the average of the preceding, and so is the case moving one additional month backward. Going all the way back to month 1 we have just one single observation 2 and then “expect” 10 altogether for the 5 months ahead, and that is exactly what we got, thus leading to a high tail probability.

Exercises

1. Replace the number 4 of the sixth month with 3 and repeat the analysis.
2. Analyse the data available after five months, i.e. omit month 6.
3. Use the described method to analyse the second example

3.5 Explaining adverse events: Categorical regression

We want to investigate how some combinations of the variables in $X = (X_1, X_2, ..., X_r)$ may trigger, predict or explain an adverse event, here denoted by Y=1 if the event occurs and 0 otherwise. The explanatory variables in X may be of different kinds: numerical and/or categorical.

This problem is common to many fields ranging from banking (credit scoring) to medicine (response studies, survival studies), for which theory is well developed and widely applied. There are two different types of studies: (i) Follow-up studies, where the individual with a given X is followed until the result Y is materialized, (ii) case-control studies, where we first observe Y, and then try to relate It to the "background" X. Reliable inferences are easiest to obtain for the first type (prospective), while the second type (retrospective) requires greater care to avoid misinterpretation of data.

For risk management in an enterprise, we may think of two different contexts as well:

In system planning: We try out different combinations of input variables, and observe whether they lead to a predefined adverse event or not.

In investigation: An adverse event of a given kind has occurred repeatedly, and we pick up the background information for each event, and at the same time we collect sufficient background information for situations not leading to the adverse event.

A fruitful line of thinking is to imagine a linear score function

$$Z = \beta_0 + \beta_1 \cdot X_1 + \beta_2 \cdot X_2 + ... + \beta_r \cdot X_r$$

and that the adverse event occurs when the score Z exceed a threshold T, i.e.

$$\begin{aligned} Y &= 1 \quad if\ Z \geq T \\ &= 0 \quad if\ Z < T \end{aligned}$$

If we think of this threshold as a random variable with cumulative distribution function $F(t) = P(T \leq t)$ (non-decreasing from 0 to 1), we get for given $Z = z$

$$P(Y = 1 | Z = z) = F(z)$$

For given $X = x$ where $x = (x_1, x_2, ..., x_r)$ we then have

$$P(Y = 1 | X = x) = F(\beta_0 + \beta_1 \cdot x_1 + \beta_2 \cdot x_2 + ... + \beta_r \cdot x_r)$$

The function F is called the link –function, and a suitable choice is the logistic function

$$F(z) = L(z) = \frac{e^z}{1+e^z}$$

Data analysis based on the logistic function is named *logit* analysis. The alternative of using the cumulative standard normal distribution G(·) instead , is named probit-analysis. Other choices of F(·) exist, but in practice the choice plays a minor role, and logit is preferred as long as the data does not tell you otherwise.

A categorical regression analysis may involve explanatory variables that are numerical and/or categorical (dichotome, nominal, ordinal). Categorical variables may be represented numerically by 0-1 variables (also named indicators or dummies). A category variable with k categories requires k-1 indicators to represent the k categories, taking one category as basis with 0 code for all k-1 variables. In practice you do not have to do this coding, since statistical software typically offers direct specification of category explanatory variables A, B, …(often named factors). Several factors may possibly interact, and they may possess a specific structure ”crossed” or ”nested”, and software may offer the opportunity to specify this.

Consider the case where X is a scalar, i.e. r=1. We then have

$$P(Y=1 \mid X=x) = \frac{e^{\beta_0+\beta_1\cdot x}}{1+e^{\beta_0+\beta_1\cdot x}}$$

The ratio between the probabilities for Y=1 and Y=0 for a given X=x is named ”the odds” for given X=x. For a logistic model this is given by

$$O = \frac{P(Y=1 \mid X=x)}{P(Y=0 \mid X=x)} = e^{\beta_0+\beta_1\cdot x}$$

The natural logarithm so-called ”*log-odds*” or *logit*, is then

$$\log\frac{P(Y=1 \mid X=x)}{P(Y=0 \mid X=x)} = \beta_0 + \beta_1 \cdot x$$

i.e. a linear function of x .

In the logistic model β_1 may be interpreted as the change in log-odds by changing X by one unit, i.e. like the interpretation of a regression coefficient. The odds itself is then changed by the multiplicative factor e^{β_1}. Then $\beta_1 = 0$ corresponds to the probability of Y=1 not depending on X at all. Then the (log-)odds is constant equal to β_0, where $\beta_0 = 0$ corresponds to the probability of Y=1 and Y=0 both being ½.

A special case of interest is when X takes only two values, 0 and 1, representing two risk groups, those who are exposed to a specific risk factor (X=1) and those who are not (X=0). We then have

$$\frac{P(Y=1 \mid X=0)}{P(Y=0 \mid X=0)} = e^{\beta_0} \text{ and } \frac{P(Y=1 \mid X=1)}{P(Y=0 \mid X=1)} = e^{\beta_0+\beta_1}$$

It is also of interest to compare the probabilities of the adverse event for the two risk groups by the so-called risk-ratio (RR)

$$RR = \frac{P(Y=1 \mid X=1)}{P(Y=1 \mid X=0)}$$

The expression becomes a bit more complicated and is omitted here (write it down yourself!). The uncritical use of this concept is pointed out elsewhere in these notes.

In theory and practice it is convenient to have the concept of ”odds-ratio”:

$$OR = \frac{P(Y=1 \mid X=1)/P(Y=0 \mid X=1)}{P(Y=1 \mid X=0)/P(Y=0 \mid X=0)} = e^{\beta_1}$$

In the case that X does not influence Y we have $\beta_1 = 0$, and therefore OR=1. In the case that X is a numerical variable, the odds-ration describes how the odds is changed when the variable X is changed by one unit.

It may be illuminating to consider the 2x2 situation from the probability table of the unconditional probabilities, where p_{ij} denotes P(X=i,Y=j):

	Y=0	Y=1	Sum
X=0	p_{00}	p_{01}	$p_{00}+p_{01}$
X=1	p_{10}	p_{11}	$p_{00}+p_{01}$
Sum	$p_{00}+p_{10}$	$p_{01}+p_{11}$	1

By computing the conditional probabilities and substitution in the OR formula we get

$$OR = \frac{p_{00}\,p_{11}}{p_{10}\,p_{01}}$$

From this we see that X and Y appears symmetric in the formula, so that we formally get the same result if exchange Y and X in the definition of OR. This does not make much sense when we think of risk attached to Y for given X. However, this throws some light on what is possible to achieve from data. With a prospective follow-up study, with preselected exposure groups we face no problems. In the case-control situation, where the adverse events are observed first, and then the circumstance is revealed, we may face a problem.

We see that OR is uniquely determined by β_1. According to the symmetry, it is possible to estimate β_1 even if the data is obtained in the reverse order X for given Y. To predict Y for given X, we also need an estimate of the constant term β_0. In some cases it is possible to get his by other means. Example: In epidemiology we may have the fraction of deaths and alive in the population for which we want to predict the effect of an additive, where part of the population is exposed and the rest not. These considerations also hold in the case of nominal and scale variables, and to several explanatory variables. This is a nice property of the logistic model, not shared by others.

Estimation of the parameters in a logistic regression model is typically performed by the ”maximum likelihood” principle, which amounts to finding the parameters that makes the observed result most probable. For the logistic model there is no simple formula, as we have for standard linear regression with independent, identically distributed normally error terms, (which then coincides with the common least squares estimate). The calculations are therefore performed by a maximization algorithm on a computer, which reaches the maximum by iterations.

For the case where X is a vector, the theory above may be extended, but left out here.

Example Challenger disaster

On January 28, 1986 millions of TV viewers watched the launch of the space shuttle Challenger from the launch site in Florida. The launch was seemingly perfect, but then 73 seconds into the flight the rocket booster seemingly exploded and killed all seven of the crew, among them the first civilian in space, a schoolteacher. The investigation commission appointed by President Reagan found that the external fuel tank had collapsed and that the likely cause of this was a burn through of an O-ring seal at a joint in one of the solid-fuel rocket boosters. After each launch the boosters were recovered and inspected. Of the previous 24 shuttle launches, 7 had damage to the joints, 16 had no damage and 1 was unknown because the boosters were not recovered. The commission did not settle at any cause for the damage, perhaps because they did not compare with the no damage cases. Later it was pointed out that January 28, 1986 was a particularly cold day at the launch site, about 15 °F cooler than any previous launch, and the question was raised: Could low temperature contribute to the accident?

The temperature data for the 23 launches with available information was as follows:

Temp ˚F	Damage	Temp ˚F	Damage	Temp ˚F	Damage
66	0	57	1	70	0
70	1	63	1	81	0
69	0	70	1	76	0
68	0	78	0	79	0
67	0	67	0	75	1
72	0	53	1	76	0
73	0	67	0	58	1
70	0	75	0		

Here is a dotplot showing that the cases with damage are more present at the lower temperatures.

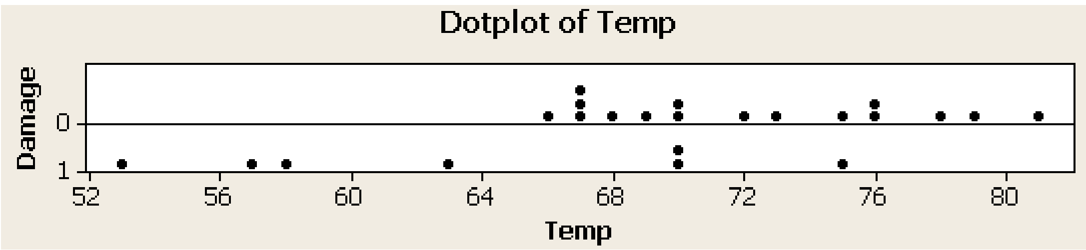

If we model the probability of damage as function of the temperature by the logistic model, and run a logistic regression, we get the following output. We see that temperature is statistically significant at the 5% significance level (P-value 3.2%). The odds-ratio is 0.79, which means that the odds for having damaged joints are just reduced to 79% of its size by increasing the launching temperature by one degree Fahrenheit.

Binary Logistic Regression: Damage versus Temp

```
Link Function: Logit

Variable  Value  Count
Damage    1          7  (Event)
          0         16
          Total     23

Logistic Regression Table
                                                  Odds     95% CI
Predictor       Coef   SE Coef      Z       P  Ratio  Lower  Upper
Constant     15,0429   7,37862   2,04   0,041
Temp       -0,232163  0,108236  -2,14   0,032   0,79   0,64   0,98

Log-Likelihood = -10,158
Test that all slopes are zero: G = 7,952, DF = 1, P-Value = 0,005

Goodness-of-Fit Tests

Method           Chi-Square  DF      P
Pearson             11,1303  14  0,676
Deviance            11,9974  14  0,607
Hosmer-Lemeshow      9,7119   8  0,286
```

The probabilities of damaged joints at some different temperatures, calculated from the estimated model are given in the table and graphed below. We see that with temperatures down in the 40's we are pretty sure to get in trouble!

Temp	50	55	60	65	70	75	80	85	90
Prob	0.97	0.91	0.75	0.49	0.23	0.09	0.03	0.01	0.003

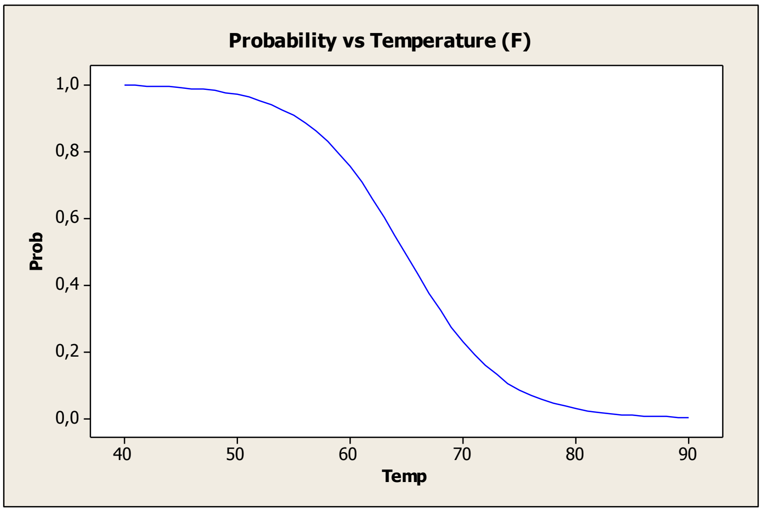

Exercise

The commission for the Challenger accident presented the following graph in their report, showing the number of damaged joints for each of seven launches with damage plotted against, the temperature at launch.

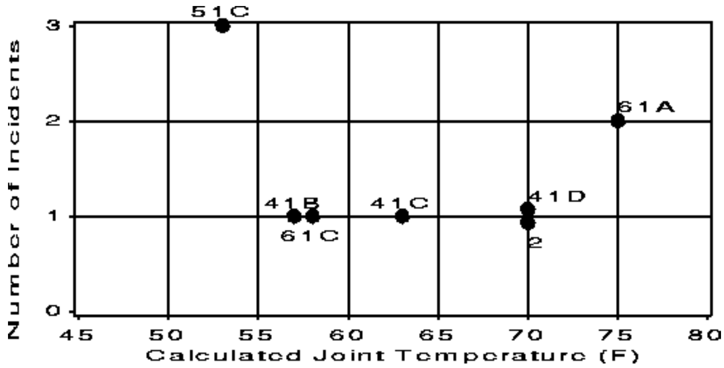

Comment the graph, and a possible false conclusion derived from the graph.

We close this session by presenting an important theoretical property of the logistic model, with tells that it is useful in practice also for retrospective studies. Recall that

<u>Prospective study:</u> Take objects with given X's and observe Y's
<u>Retrospective study</u>: Take objects from each of Y=1 and Y=0 and observe their X's

Assume as before the logit model:
$$P(Y=1 \mid X=x) = \frac{e^{\beta_0+\beta_1 x}}{1+e^{\beta_0+\beta_1 x}}$$

Let the probability of picking a specific object from group no. i be q_i for i=0,1. The probability that a randomly sampled object (S) belongs to group no. i is (without knowledge of X)

$$P(Y=i \mid S) = \frac{q_i}{q_0+q_1} \quad i=0,1$$

It then follows from Bayes law that

$$P(Y=1 \mid X=x, S) = \frac{e^{\beta_0'+\beta_1 x}}{1+e^{\beta_0'+\beta_1 x}} \quad \text{where} \quad \beta_0' = \beta_0 + \log(q_1/q_0)$$

This means that sampled object still follows the logit model with the same regression coefficient. but a different constant term. This means that the effect of the explanatory variable(s) may still be examined by logistic regression, even if the objects are sampled from each of the two groups, and their explanatory variables are recorded retrospectively. Moreover if we know the sampled probabilities, we can from the constant term of this regression recover the constant term of the prospective model itself. This is a property not shared by other link functions than the logistic.

3.6 Extreme value analysis

Events with imagined low probability and large consequences are often referred to as extreme events. Issues related to extreme events are found in many areas, among others in planning for protection against environmental events (e.g. dikes for high tides) and in planning for necessary capital in order to withstand extreme financial losses (e.g. bankruptcy). Common to these areas is that we have some kind of size variable X that varies over time, in continuous time (ex. tidal height) or in discrete time (ex. daily banking loss). Common is also the need to get an idea of the probability distribution of X. How to proceed, may depend on the availability of data, whether data is abundant, scarce or not available at all. In case of scarce data, we typically have to combine data with some assumptions about the distribution, and its development over time. Knowledge of some distribution theory and some process theory may then be useful. A class of distributions of specific interest is the so-called extreme value distributions. In case of no data, we typically have to generate scenarios by simulation, and the theory mentioned may be just as useful.

For most natural disasters, the impact will increase with size or magnitude. Examples are the flood level and the magnitude of earthquakes (say measured on the Richter scale). Typically there is an inverse relationship between frequency and magnitude, so that large size events occur less frequent than small size events. A useful concept in this context is the so-called recurrence interval, also named return period, which is a quantity based on historical data, but intended use is to forecast the expected time between events or the probability of occurrence in a given time period. There are several ways of establishing a relation between the recurrence time and the magnitude based on data from sufficiently many events. Some are fairly ad hoc (but practical) and some are linked to extreme value theory. We will first consider a simple graphical technique, often (mis)named frequency analysis.

Let us imagine a river where we have observed the peak magnitude (level or flow) in consecutive years. . From these observed worst cases each year we want to say something about the risk of even worse cases in the coming years. Let $X_1, X_2,\ldots, X_n$ be the peak magnitudes for each of n consecutive years These are then ranked from the largest to the smallest $X_{(1)}\geq X_{(2)}\geq\ldots\geq X_{(n)}$,so that the m'th largest is $X_{(m)}$, The probability of a peak magnitude that exceeds $X_{(m)}$, in any given year is roughly taken to be

$$P_m=P(X\geq X_{(m)})=m/(n+1)$$

The expected time for another event of at least this size is then R= 1/P *(*often named The recurrence interval). The idea is to be able to use this for other combinations values of peaks values and probabilities as well, and hopefully also be able to extrapolate to extreme magnitudes. This is achieved by smoothing the data, for instance by plotting observed magnitudes $X_{(m)}$ against some function of the corresponding P_m (or equivalently the R's). The points in the graph may then be fitted to a smooth curve, preferably a straight line. It turns out that plotting against log(P)=-log(R) works well in many cases, but alternatives exist. We may then plot against R, but use a logarithmic scale on the horizontal axis.

Example Flood frequency

The Red River of the North runs through North Dakota into Lake Ontario in Canada. It is frequently flooded. The discharge is measured continuously at different locations and peak discharges each year are available from the US Geological Survey database, see http://www.usgs.gov/water/. Here we look at the data from Wahpeton, Richland County (Been there!), which have records from 1942 on. The maximum over the 70 years 1942-2011 is 12 800 which occurred in 1997 on April 15. The file also includes the number 100 years prior to this in 1897, not used in our analysis (left panel). In the right panel we have plotted the peak discharges in cubic feet per second against the recurrence interval R on the logarithmic scale. We see that the points are close to a straight line, and this is fitted by linear regression. By extending the line to R=100 we read a discharge of X=14000, which may be named the 100-year flood in popular terms.

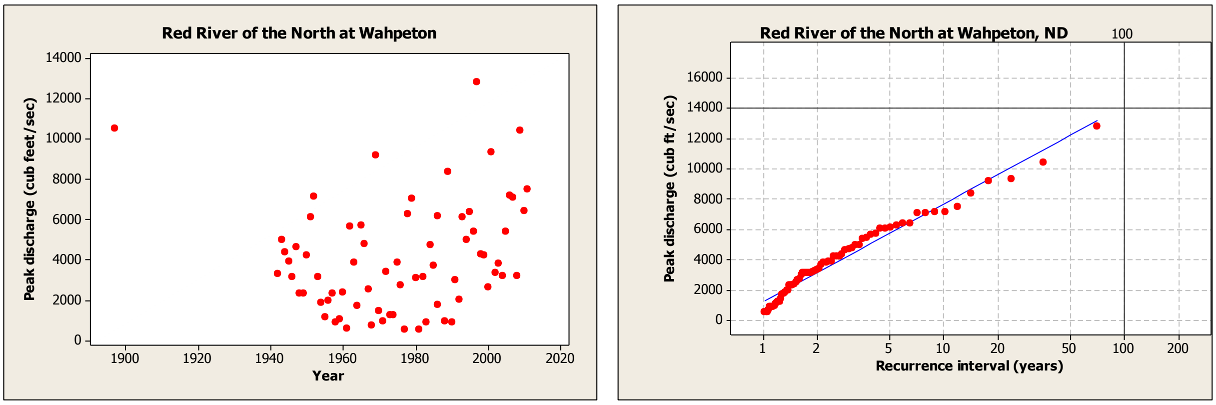

Remark. The formula above is sometimes named the Weibull-formula, and is widely used in the US. There are several alternatives, and P=(m-0.4)/(n+0.2) are more frequently used in Europe and Canada. The different alternatives are typically of form P= (m-a)/(n+1-2a) for a chosen a.

This kind of frequency analysis is mainly based on the assumptions of constancy and independence, which may be questionable, but partly possible to account for. There may also be a problem of questionable data quality, in particular since measurements are taken under extreme conditions. Anyway, the practitioners seem to be willing to extrapolate up to twice the number of years they have data, which is rather bold I would say.

More sophisticated analysis of this kind on data may be based on extreme value theory, which has found its application in many areas, whether it be natural hazards or financial risks. A key question in extreme value theory is the following

Let $X_1, X_2,\dots, X_n$ be independent and identically distributed (i.i.d.) variables with (cumulative) distribution function $F(x)=P(X\le x)$ and let $M_n = \max(X_1, X_2,\dots, X_n)$.

What is the distribution $F_n(x)$ of M_n?

Due to the i.i.d. condition we clearly have $F_n(x)=(F(x))^n$

For some classes of distributions F we have that F_n is of the same class, and it is just a change of parameters. For most distributions F the F_n does not belong to a recognizable distribution family, but is just given by the above formula.

In some cases we do not know F, and the question is whether we still can say something. In fact we can! By a suitable "normalization" with respect to n, bringing M_n on a common location and scale, our hope is that when n tends to infinity, we get a limiting distribution that also gives us approximate probabilities for large n. This is similar to more well-known situation of sum of independent and identically distributed variables, where the standardized sum has a distribution that tends to the standard normal distribution. More precisely we have:
Assume the existence of sequence of constants $\{a_n\}$ and positive sequence of constants $\{b_n\}$ so that

$$P((M_n - a_n)/ b_n \leq z) \rightarrow H(z) \text{ as } n \rightarrow \infty$$

If H is a proper distribution, it has to be one of three types:

Frechet: $$H(z) = \begin{cases} 0 & ; \ z < a \\ \exp\left\{-\left(\frac{z-a}{b}\right)^{-\alpha}\right\} & ; \ z \geq a \end{cases}$$

Weibull: $$H(z) = \begin{cases} \exp\left\{-(-\left(\frac{z-a}{b}\right))^{-\alpha}\right\} & ; \ z < a \\ 1 & ; \ z \geq a \end{cases}$$

Gumbel: $$H(z) = \exp\left\{-exp(-\left(\frac{z-a}{b}\right))\right\}$$

These three classes of distributions are named the extreme value distributions. We note that if the observations come from any of these distributions at the outset, their maximum will follow a distribution of the same kind with just a change of parameters. Moreover, if the observations come from a different distribution, their maximum may nevertheless approximately follow an extreme value distribution. Which one will depend on its corresponding "domain of attraction". This seem to imply that, under the maximum of i.i.d. variables, and when the main concern is extreme tail of the distribution, there is no need to consider distributions outside the classes of extreme value distributions. However, there may be several reasons for other choices: Some want to have a good description of the central part of the distribution as well, and some prefer distributions, which better provides a consistent and meaningful modeling approach tailored to their specific field.

Example:
Consider two variables X and Y and 10 observations of each, here recorded in increasing order.

X:	10	13	16	19	21	28	32	35	38	43
Y:	14	17	18	20	22	23	25	26	38	52

Descriptive Statistics: X; Y

Variable	Mean	StDev	Minimum	Q1	Median	Q3	Maximum
X	25.50	11.31	10.00	15.25	24.50	35.75	43.00
Y	25.50	11.39	14.00	17.75	22.50	29.00	52.00

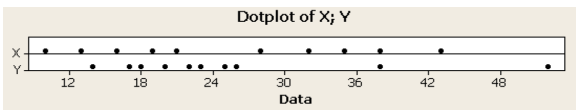

The descriptive statistics show that they have the same mean and (approximately) the same standard deviation. Fitting a normal distribution with these characteristics to the data and calculation of the risk of surpassing high levels, say 50 and 60, may be greatly flawed. We see from the dotplot that the distributions of values are quite different, and both far from normal: X may correspond to an even distribution within a range, while Y has a clear outlying observation. This is scarce data, and more data may reveal observations far beyond those obtained so far. For Y we are already warned, but for X we may not be prepared for surprises. In cases like this it is clearly of value to have additional information, in particular in the case of X. Say, we know that the possible range of values, the data will support a uniform distribution over this range. If we do not know the range, and are worried about extremes, we may judge the value of further observation. The overall fit to a particular distribution may be judged by so-called probability plots, in which the points should be approximately on a straight line to justify the specific distribution the plot is designed for. Below follows plots for the Y-data for the normal, lognormal and two distributions with additional shift parameter. We see that the normal is clearly rejected, the lognormal is fairly good, and improved by adding the shift parameter. Corresponding plot for the extreme value distributions will show that they do not give a good all over fit, but quite often it is enough to have good fit in the upper tail, so this does not rule out the rationale for using them for risk calculation.

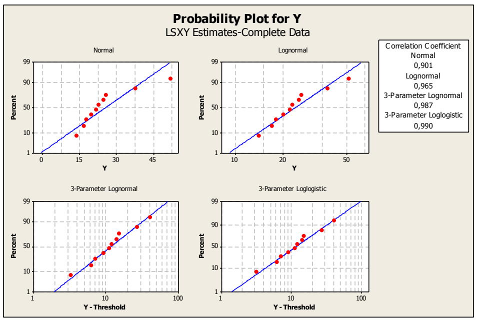

We will take a closer look at one of the extreme value distributions, namely the Frechet distribution, which has found application in many areas, among then hydrology (flood control) and finance. The distribution function of the standard form (a=0, b=1) is given by

$$F(x) = P(X \leq x) = e^{-x^{-\alpha}} \quad \text{for } x \geq 0$$

This means that the upper tail probability is

$$P(X > x) = 1 - e^{-x^{-\alpha}} \approx x^{-\alpha}$$

which is decreasing exponentially with α. It turns out that all moments of order < α exist (and no moment of order $\geq \alpha$). Thus the expectation does not exist when $\alpha \leq 1$.
The graph illustrates the distribution densities for α=0.5, 1 and 2 (peaks left to right)

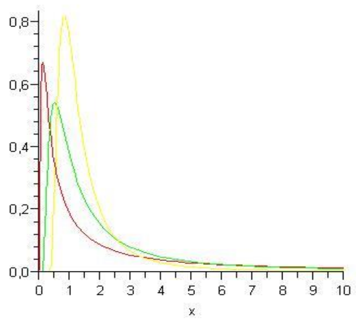

Here are histograms for 100 simulated observations from the Frechet distribution α=1.5, 1 and 0.5

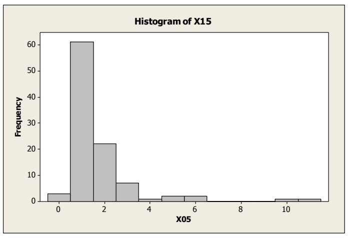

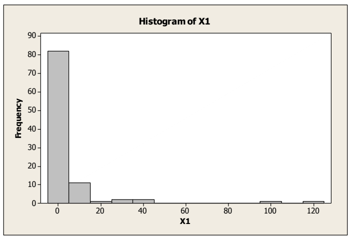

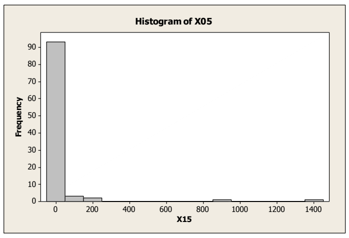

The upper p-quantile of the standard Frechet distribution is given by

$$x_p = (-\ln p)^{-1/\alpha}$$

Here are the upper 99% quantiles for some choices of α:

α	0.5	1.0	2.0	3.0	4.0
99% quantile	9900.08	99.50	9.97	4.63	3.16

The fitting of distribution and calculating quantiles may in principle go as follows:

1. Estimate the tail index α
2. Estimate the distribution (location and scale) with the tail index estimate given
3. Estimate by computing the p-quantile with the estimated distribution

The tails of other heavy tailed distributions, like student-t, Pareto and logGamma, typically decrease exponentially the same way, and the corresponding α may therefore be used to characterize the heaviness of the upper tail in general. The above shows that information about the tail-index may be vital for assessing risk, as well as for planning, e.g. the height of dikes, the capital needed to withstand extreme losses.

$$Q(x) = P(X > x) \sim kx^{-\alpha} \quad \text{for large x}$$

Then

$$lnQ(x) \sim -\alpha ln x + \beta$$

This provides an opportunity for graphical identification and estimation.
Take $x_1 \leq x_2 \leq \cdots \leq x_n$ and let Q_i= No. of observations $>x_i$.

Plot lnQ_i against lnx_i and look for tail linearity. If so, draw the line and read off α and β, and take $k = e^{\beta}$.

<u>Example:</u> A case of n=100 independent observations gave the following graph, where we see a linear trend for high values with slope about -1 , corresponding to α=1.

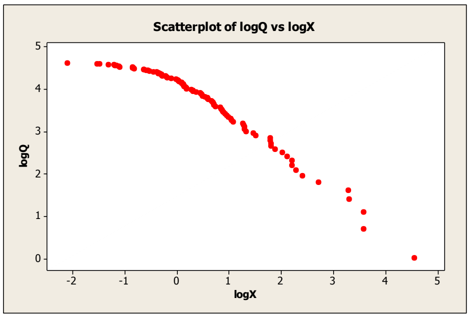

It is also of interest to look at extreme values, given that the value is above a given threshold t, often named Peaks Over Threshold (POT). The probability of interest is therefore

$$P(X > x \mid X > t) \text{ for } x > t$$

There are several reasons for this:

1. Only values above a threshold are reported
2. It may provide alternative way of assessing the risk of extremes

As for the second reason note that for x>t

$$P(X > x) = P(X > t) \cdot P(X > x \mid X > t)$$

If we can say something about both factors on the right hand side for high t, where we still have data, we may be able to say something about the tail probability for even higher x as well, in the region where we have no data. Extreme value theory again comes to help, saying that the second factor may be evaluated on basis of the so-called Generalized Pareto Distribution (GPD)
What have we obtained?

- Assessment of the risk of surpassing extreme values that have never occurred before.
- Assessment of extreme levels not being surpassed with a given probability assurance.

But......
"There is always going to be an element of doubt, as one is extrapolating into areas one doesn't know about. But what EVT is doing is making the best use of whatever data you have about extreme phenomena" (Richard L. Smith)

3.7 Survival models and processes

Waiting time distributions and hazards

Let T be the waiting time from time zero until a specific event (e.g. an accident or a fatality). Then the cumulative probability distribution is given by

$$F(t) = P(T \le t)$$

$S(t) = P(T > t) = 1 - F(t)$ is named the survival function of T.

Assuming a continuous distribution we can represent by a probability density $f(t)$ so that

$$F(t) = \int_{-\infty}^{t} f(u)du$$

We see that the density is the derivative of the cumulative distribution function, i.e. $f(t) = F^{'}(t)$.
$f(t)dt$ may be interpreted as the probability of an event in the period $(t, t+dt)$.
The hazard-rate function (or failure rate) is given by

$$r(t) = \frac{f(t)}{S(t)} = \frac{f(t)}{1-F(t)}$$

Here $r(t)dt$ is the probability of an event in the period $(t, t+dt)$ given no events in [0, t].

The expectation of T is given by (often denoted by MTTF "Mean Time To Failure")

$$ET = \int_{-\infty}^{\infty} t \cdot f(t)dt$$

with no contribution to the integral for negative t.

Specific distributions

T is distributed exponentially with parameter $\lambda > 0$ if

$$F(t) = 1 - e^{-\lambda t} \quad t \ge 0$$

and so

$$f(t) = \lambda e^{-\lambda t} \quad t \ge 0$$

In this case $r(t) = \lambda$ i.e. constant hazard rate and

$$P(T > u+t \mid T > u) = \frac{P(T > u+t)}{P(T > u)} = \frac{e^{-\lambda(u+t)}}{e^{-\lambda u}} = e^{-\lambda t}$$ i.e. memoryless!

The expectation of T becomes

$$ET = \int_0^\infty t \cdot \lambda e^{-\lambda t} dt = \frac{1}{\lambda}$$

T is said to have a Weibull distribution if its cumulative distribution is given by

$$F(t) = 1 - e^{-(\lambda t)^\alpha} \qquad t \geq 0$$

where $\lambda > 0$ and $\alpha > 0$ are parameters that characterize the distribution.
The hazard-rate function in this case is

$$r(t) = \alpha\lambda(\lambda t)^{\alpha-1}$$

so that we have increasing hazard rate for $\alpha > 1$ and decreasing hazard rate for for $\alpha < 1$

Components that wear out will have increasing failure rate (IFR). Many electronic components seem to have close to constant failure rate, and this is often taken as a simplifying assumption. Some systems that are continuously improved during operation may have decreasing failure rate (DFR). However, many systems typically have decreasing failure rate at the beginning, when some flaws are removed, then fairly constant failure rate for an operational period, and then increasing failure rate as the system wears out despite maintenance. In such cases the failure rate function has appearance like a "bathtub".

Proportional hazard models

In many situations we want to model the relationship between the waiting times and one or more explanatory variables. One way of doing that is by *the proportional hazard model*, which assumes that the effect of an increase of one unit in a given variable is multiplicative with respect to the hazard rate. The hazard rate for given $X_1, X_2, \dots, X_r$ may be written as

$$r(t) = r_o(t) \cdot e^{\beta_1 X_1 + \beta_2 X_2 + \beta_r X_r}$$

where $r_o(t)$ is a baseline hazard rate function. Then e^β is the multiplicative factor to the hazed rate if we change the variable X by one unit to X+1. The explanatory variables may be numerical or 0-1 variables like in ordinary regression. In case of category variables with more than two categories, the categories have to be represented by 0-1 variables (often named indicators or dummies). A category variable with k categories is represented by k-1 indicators, with a base category having 0 code on all.

<u>Example</u>: In a study of survival of patients explanatory variables may be such as age, sosio-demographic variables, type of treatment and dose. Age may have to be categorized (why?).

This model has found wide application in many areas and with user-friendly software the interpretation of output is quite similar to that of ordinary regression.

Competing risk

Consider a situation where we observe time to failure and the cause of failure among a given set of potential causes. Examples of this may be

Machinery: Time to failure due to wear or negligence
Humans: Life length and registered cause of death

In many contexts we can imagine that the causes are "competing", and that we are interested in understanding how the failure causes contribute to time to failure.
What we observe for a given case may be considered as random, both with respect to time to failure and cause. We therefore consider a random pair (T,C) where

T = Time to failure (or Life length)
C = Cause of failure (one of numbers 1,2,…,m)

and (T,C) having a joint probability distribution that reflects the interplay of the possible causes. We then assume n observations sampled from this distribution are the basis for inference about this distribution.

In order to model the competing risk aspect it may be useful to imagine a life length T_i attributed to each cause i, which is the life length if this cause was the only one in effect. This life is unobservable (often named latent), but what we observe is T = minimum(T_i).

$F_i(t) = P(T_i \le t)$	Cumulative distributions
$S_i(t) = P(T_i > t)$	Survival functions
$h_i(t) = f_i(t) / S_i(t)$	Hazard rates

We then have for T=minimum(T_i)

$$S(t) = P(T > t) = P(all\ T_i > t)$$

which may be complicated in general, but in case of the T_i's independent we have

$$S(t) = P(\bigcap_{i=1}^{m}(T_i > t)) = \prod_{i=1}^{m} P(T_i > t) = \prod_{i=1}^{m} S_i(t)$$

In general the hazard rate for T is

$$h(t) = f(t) / S(t)$$

The statistical problem of making inference on the separate causes from observations of T is a challenging one, with an extensive literature.

Some examples of applications for independent exponentially distributed waiting times:

<u>Example:</u> Failure due to hits by independent causes
Suppose that each of m causes of failure has exponential waiting times until it hits, with hit rates λ_i, i=1,2,…,m. The life length T of the system (i.e. the time until the first cause hits) is then also exponentially distributed with rate sum of the λ_i's. For equal $\lambda_i = \lambda$ exponential with rate mλ.

<u>Example</u>: Failure of systems in series and parallel
The same argument can be used for a system with components in series which works as long as all components are working, i.e. the life length T is the minimum of the life length of the components. For a system with redundant components in parallel the life length of the system is the maximum of the life length of the components, and this distribution is more complicated, but the expectation is easily derived using the memoryless property of the exponential distribution. For m components with equal $\lambda_i = \lambda$ we have

$$ET = \frac{1}{\lambda}\left(1 + \frac{1}{2} + \frac{1}{3} + \cdots + \frac{1}{m}\right)$$

Argument: Expected waiting time until the first of the m components fails is 1/mλ. On this failure there are m-1 components under risk and the expected time until another fails is 1/(m-1)λ regardless how long we waited for the first failure etc. until one component remains with expected life 1/λ.

<u>Example:</u> Cascading failure of parallel system
Consider a redundant parallel system with two components working independently with exponential lifetime with failure rate λ. When the first component fails the other experiences increased stress for a period of length δ. The expected lifetime of the system turns out to be

$$ET = \frac{1}{\lambda} + \frac{1}{2\lambda} e^{-2\lambda\delta}$$

This is decreasing in δ and note that $\frac{3}{2\lambda} < ET < \frac{1}{\lambda}$, the bounds corresponding to the special cases δ=0 (no increased stress) and $\delta=\infty$ (permanent increased stress)

The Poisson Process:

The number of events N_t up to time t is Poisson distributed with expectation λt , i.e.

$$P(N_t = n) = \frac{(\lambda t)^n}{n!} e^{-\lambda t} \quad n = 0,1,2,...$$

This may be derived under the assumption of events occurring randomly at constant rate λ per time unit independent of each other.

Note that $P(T > t) = P(N_t = 0) = e^{-\lambda t}$, so that the waiting time to the first event is exponentially distributed with parameter λ and expectation $1/\lambda$. It follows also that the waiting times between successive events, are all exponentially distributed with parameter λ and independent of each other, and this may be used to simulate the process. The waiting time to the n'th event T_n will then have a Gamma(n,λ) distribution.

<u>Example</u>: N_t the number of accidents up to time t.
Let p be the probability that an accident leads to a fatality, and assume that this happens independent of prior events. Then the number of fatalities up to time t is still a Poisson process, but with λ replaced by λp . In general, such a derived process is called a thinned Poisson process.

<u>Example</u>: Warning system for rocket attack

Model: Assume

1. False alarms come according to a Poisson process with rate λ, say 200/year
2. Time to decide if alarm is false is constant equal to r, say r=4 min.

Worry: New alarm arrives while the previous is still under examination.

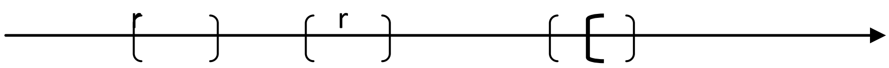

Let W=time to arrival of next alarm, while another is under examination
Problem: Find EW, and discuss its dependence on r and of λ.
Solution:
An alarm not detected false before a new alarm means an arrival before time r, and the probability of this is $p = 1 - e^{-r\lambda}$.
Let N=the number of alarms that is cleared before the first one arrives while still examining an alarm. N is obviously geometrically distributed, so that $E(N) = 1/p$. We then get by double expectation

$$E(W) = E(E(W|N) = E\left(\frac{N}{\lambda}\right) = \frac{1}{\lambda} \cdot \frac{1}{1 - e^{-r\lambda}} \approx \frac{1}{\lambda} \cdot \frac{1}{r\lambda} = \frac{1}{r\lambda^2}$$

Lesson: If you double the arrival rate, the expected time to the first critical situation is reduced to 1 / 4 !

With 200 false alarms a year and 4 minutes to check if false we get with year as time unit λ=200 and $r = 4/(60 \cdot 24 \cdot 364)$ and so $E(W) = 3.276$. So, if the arrival rate is doubled, the expected time to criticality will be less than a year.

The Compound Poisson Process:

A process which can be represented as

$$Y_t = X_1 + X_2 + \ldots + X_{N_t}$$

where N_t is a Poisson process and $\{X_t, i = 1, 2, ..\}$ is a sequence of independent identically distributes variables and independent of N_t.

$$E(Y_t) = \lambda t \cdot E(X_i) \qquad \text{var}(Y_t) = \lambda t \cdot E(X_i^2)$$

<u>Example:</u> Y_t = sum of costs associated with each event occurring up to time t.
Going back to the previous example, we may think of initiating events (accidents), which with probability p lead to fatality costs. The sum of the fatality cost will the be a compound Poisson process as well, with λ replaced by λp.

3.8 Risk simulation

Simulation is used in practice when a theoretical model does not provide an exact analytical solution. An exact solution may possibly be obtained by making simplifying, but questionable assumptions. Simulation allows both more realistic models and more complicated models, and also provides the opportunity to vary the assumptions in order to study the effect on the solution. When demanding mathematics appears as a barrier, simulation may also come to rescue. Among the many advantages simulation may have over an analytical solution are

- a low threshold for analyst, with more
- easily communicated model descriptions, assumptions and results to stakeholders

Simulation is nowadays applied in many areas due to almost unlimited computer capacity and user-friendly software, among others in operations management, investment analysis and risk analysis. The opportunities are many, among them:

- describe a system in order to study its features on the computer
- generate data for decision support, e.g. for optimization
- generate scenarios under different assumptions

Simulation is ideal for the study of issues containing randomness and uncertainty, and we then use the term stochastic simulation or Monte Carlo simulation is quite often used. This is based on stochastic models, and performed by random drawings of the random variables involved, and then compute the value of some criterion variable. This is repeated many times, in order to get at the probability distribution of the criterion variable and thereby reveal the risk and opportunity involved.

Example 1 Project cost
Consider a housing project with five cost elements, each characterized by three outcomes: Minimum, Most likely (“best guess”) and Maximum cost, as follows

	Minimum	Best guess	Maximum
Site	200	300	500
Foundation	300	400	500
Building	600	800	1200
Interior	500	700	800
Exterior	200	300	400

If these numbers are used in a ”if so, then” analysis there are 3^5 = 243 possible combinations for a specific scenario. This is too many to be of any use separately. It is not wise to compute the cost for all the 243 combinations and then compute the average. By doing this we neglect that the costs can be any number between the given minimum and maximum, and also neglect that the best guess is more likely than the extremes.

It is more reasonable to imagine that each of the five have a probability distribution with range from the minimum to the maximum. In case take the maximum as if it has a 1% chance of being reached, the probability that all the costs reach the maximum is 0.01^5 which is 1 : 10 000 000 000. This is based on the assumption that the five cost elements are independent. In a simulation we sample outcomes from the assumed probability distribution for each of the five cost elements and add them.

When repeating these five drawings and taking their sum a large number of times, say 1000, we will have a good picture of the probability distribution of the total cost, and so the chance that the total cost will not surpass given levels. Here is the result from such a simulation illustrated graphically:

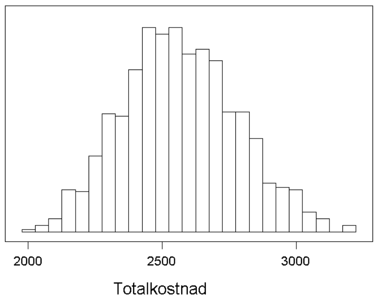

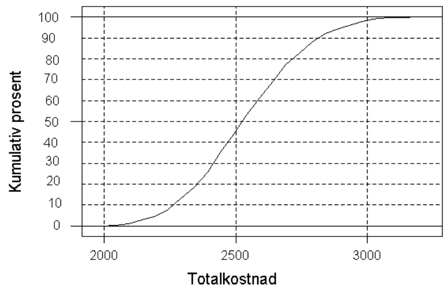

The histogram (left) gives a picture of the uncertainty in the total costs, and from the cumulative distribution (right) we can read an estimate of the probability that the total cost does not surpass different given levels. The simulations are performed under the assumption that the individual costs follow a triangular distribution. This is a simple convenient assumption, able to pick up the main feature. We return to this below.

Central to simulation is a generator of random numbers between 0 and 1. A sequence of numbers generated this way are taken to be independent drawings from the Uniform[0,1] distribution, i.e. the uniform (rectangular) distribution over the interval [0,1].

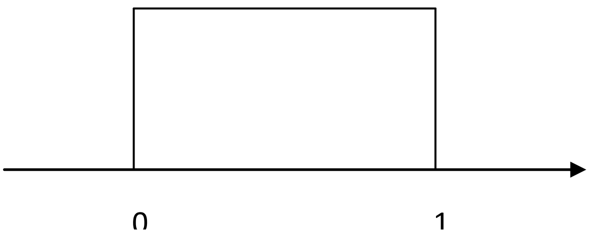

Such drawings are the basis for samples from other distributions. Special software is designed for extensive simulations and analysis that may be helpful to

- describe complicated processes
- simulate data from many different distributions
- perform diagnostic checking
- keep the information from the repeats and generate relevant and readable reports

However, the possibilities are good with standard statistical software as well, and even with spreadsheet programs like Excel. If you know the links between the uniform distribution and other distributions, it is not difficult to program simple simulations in spreadsheets.

Example : We wish to simulate a variable X with probability distribution given by

x	1	2	3	4
P(X=x)	0.3	0.4	0.2	0.1

Let U = u be simulated value from the Uniform(0,1)-distribution. We may the assign values to X as follows:

$$\begin{array}{ll} 0 \le u \le 0.3 & \to 1 \\ 0.3 \le u \le 0.7 & \to 2 \\ 0.7 \le u \le 0.9 & \to 3 \\ 0.9 \le u \le 1.0 & \to 4 \end{array}$$

This is easy to program, but with many outcomes it will become tedious to list all possible values and their probabilities. In software you are likely to find commands for important and well known discrete distributions, like the discrete uniform, binomial and Poisson.

For simulation in practice, we also need to sample from various continuous distributions. We have three kinds according to their support:

Bounded interval [a,b]:	Uniform, Triangular, Beta
Positive real line [0, ∞] :	Exponential, Gamma, Weibull, Frechet, Pareto etc.
Real line [- ∞, ∞]:	Normal, Logistic, Heavy tail distributions

The triangular distribution used in our previous example is a continuous distribution representing increasing probabilities from a smallest possible value (a) to the most likely value (c) and then diminishing probabilities to the maximum possible value (b) as shown in the graph

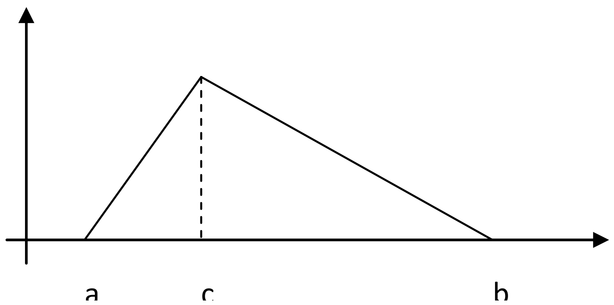

It turns out that simulation from the continuous distributions above may be obtained by fairly simple transformations of Uniform[0,1]- variables. Let us see why and how.

Example X distributed exponential with parameter λ, i.e.

$$F(x) = 1 - e^{-\lambda x} \qquad x \geq 0.$$

We can then write for U Uniform[0, 1].

$$U = F(X) = 1 - e^{-\lambda X} \Leftrightarrow X = F^{-1}(U) = -\frac{1}{\lambda}\ln(1-U)$$

Since 1 - U is also Uniform[0, 1], we can alternatively simulate X by

$$X = -\frac{1}{\lambda}\ln U$$

Example X distributed triangular(a,c,b)

The cumulative distribution is then given by

$$\begin{aligned} F(x) &= \frac{(x-a)^2}{(b-a)(c-a)} && x < c \\ &= 1 - \frac{(b-x)^2}{(b-a)(b-c)} && x \geq c \end{aligned}$$

Solving the equation U=F(X) we get

$$\begin{aligned} X &= a + \sqrt{(b-a)(c-a)U} && \textit{for} \quad U < (c-a)/(b-a) \\ &= b - \sqrt{(b-a)(b-c)(1-U)} && \textit{for} \quad U \geq (c-a)/(b-a) \end{aligned}$$

Example X distributed normal $N(\mu, \sigma^2)$

In the special case of the standard Gaussian distribution G we have that $X = G^{-1}(U)$ is Normal(0,1). However, since there is no simple expression for the inverse of the Gaussian, standard normal variates are usually generated by other methods. One possibility is to use the following amazing result:

If U_1 and U_2 are independent Uniform[0,1] we get two independent N(0,1) variables Z_1 and Z_2 by

$$\begin{aligned} Z_1 &= \sqrt{-2\log U_1}\cos(2\pi U_2) \\ Z_2 &= \sqrt{-2\log U_1}\sin(2\pi U_2) \end{aligned}$$

We then get two independent $N(\mu, \sigma^2)$ variables by

$$\begin{aligned} X_1 &= \mu + \sigma Z_1 \\ X_2 &= \mu + \sigma Z_2 \end{aligned}$$

Let X be a random variable with cumulative distribution

$$F(x) = P(X \le x)$$

Properties: F(x) is non-decreasing, like a staircase for discrete distributions and continuous for continuous distributions.

Theorem

1. For X with cumulative distribution F(x), Then U = F(X) distributed Uniform[0, 1].
2. If U is Uniform[0, 1] and F is any continuous cumulative distribution with inverse F^{-1}, then $X = F^{-1}(U)$ has cumulative distribution F.

This theorem may be the basis for simulating X from U, at least in cases where F^{-1} has a simple analytic expression.

Exercise

Show how to simulate variables from the following distributions

(i) Weibull $F(x) = 1 - e^{-(\lambda x)^{\alpha}}$

(ii) Frechet $F(x) = \exp(-x^{-\alpha})$

(iii) Pareto $F(x) = 1 - (\frac{x}{\beta})^{-\alpha} \quad x \ge \beta$

Modeling dependencies

The simulation of two or more dependent random variables X_1 and X_2 may be achieved by different means. We have mainly two situations:

1. One of the variables, say X_2, is a function of the other plus a random "error".
2. The two variables are genuinely bivariate

The first case is simple if we know the function f and the probability distribution of the error. The second case is also simple for bivariate normal variables, so let us consider the simulation of $X_1 \sim N(\mu_1, \sigma_1)$ and $X_2 \sim N(\mu_2, \sigma_2)$, where X_1 and X_2 are dependent with correlation ρ. We may then depart from two independent N(0,1) variables Z_1 and Z_0. From these we compute

$$Z_2 = \rho Z_1 + \sqrt{1-\rho^2}\ Z_0$$

Z_1 and Z_2 are then N(0,1) with the desired correlation ρ. We then compute

$$X_1 = \mu_1 + \sigma_1 Z_1 \qquad X_2 = \mu_2 + \sigma_2 Z_2$$

Correlated variates U_1, and U_2 with Uniform[0,1] marginals may in principle be obtained by

$$U_i = G(Z_i) \quad \text{i=1,2}$$

and correlated Y_1 and Y_2 from any other continuous marginal distribution *F* by taking

$$Y_i = F^{-1}(U_i) \quad \text{i=1,2}$$

The construction of covariate distributions with any given marginal, is studied extensively in the statistical literature using the concept of copulas. The construction above corresponds to the so-called Gaussian copula. Note that the correlations are between the intermediate normal variables, and are not easily expressed in terms of the wanted correlations between the final variables. This may be solved by trial and error. The construction above and copulas extends beyond the bivariate case, and has found wide applications in banking and finance for modelling correlated risks, e.g. default risks.

Example:
Here is an example of two simulation of from the Gaussian copula, with different marginals, standard normal on the left panel and t-distributed with 5 degrees of freedom in the right panel. Here the bivariate normal had correlation 0.5. We see that we are able to generate correlate data with heavier tails than the normal, which is found in many fields of application.

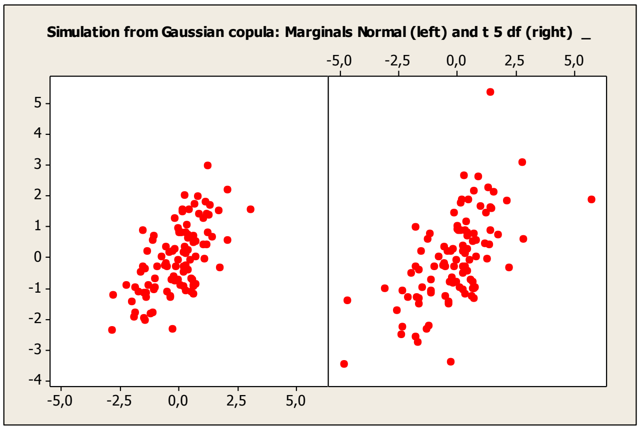

A good simulation model requires that

- the relations between the variables is well described,
- the distributions to be simulated from are realistic.

Both are a challenge, and there is a danger to overlook important aspects. In some cases the relations are simple, but dependencies are overlooked. In some cases we know that dependencies are there, but we do not know how to model them, and hope that it will not matter much if we keep them out. Example: In a project where the finishing date depends on part projects, there may be some common cause for delays. Although the description of expected time spent may represent the marginal distributions, it may not give the correct picture of the joint distributions.

The distribution of the variables involved in the simulation may be determined by

- expert judgment revealed using common principles for revelation.
- use of empirical data

and combinations of these two. For the simulation from a distribution from which we have observed data, we face three possibilities:

1. simulate directly from the empirical distribution
2. simulate from the smoothed empirical distribution
3. simulate from a fitted distribution of assumed type

We may want to remove the random character of the data by smoothing, in particular when we have little data. The fitting to a given distribution type is also a kind of smoothing, which in addition provides the opportunity to include theoretical and practical knowledge about the distribution. At the same time, we have a better regime for varying the parameters of the distribution to see the consequences of different assumptions.

Often we have the choice between a continuous and discrete distribution, and sometimes a continuous distribution be convenient even if the variable in practice is discrete. Different distribution classes are parameterized by one or more parameters, e.g. Poisson (λ), Normal (μ, σ), Gamma (α, λ), Beta (r, s). If we have n independent variables from an assumed distribution with unknown parameters, the distribution may be determined by estimating the parameters (with some confidence). A general estimation method is the "maximum likelihood" method (ML). Software exists for estimation for many different distributions, and also general algorithms for ML-estimation. Software may also provide measures of the "goodness of fit" .of the distribution to the data. This may be useful if several different distribution types are tried out for best fit. For this, we have formal statistical tests or informal graphical judgment by "probability plots". An example of this is models for elapsed time, where there are several possible models: exponential, Gamma, Weibull, extreme value etc.

Effective simulation methods

The classical form of Monte Carlo simulation is by random draws from the assumed distributions. It is desirable to determine the resulting distribution with as few repetitions as possible. Classical MC-simulation may need many repetitions to smooth out the randomness. However, quite often it is just the main feature of this distribution being of interest, and there exist different means to reduce the number of repetitions, where something is forsaken, but something else is gained. One of this is so-called Latin Hypercube Sampling (LHS). While MC-sampling may be regarded as random samples with replacement, LH-sampling may be regarded as stratified samples without replacement.

3.9 Statistics and the scientific process

Risk management needs statistics in most areas, but to a varying degree. Among others, we need statistics for measuring the frequency and severity of adverse incidences and to measure the effect of countermeasures and alternative actions. The goal is to bring forward the facts and reach a common understanding as basis for prioritizing actions. For the major issues in society, a sound scientific basis is needed. In particular, this is so for issues concerning health and environment. However, we frequently see that scientists disagree on the interpretation of the data at hand, and sometimes even disagree on the relevance of the data or the approach taken for gathering the data. This issue is imminent in connection with the so-called precautionary principle (see section 1.4).

Let us first say something about the logic of science. We have essentially three types of logical reasoning:

1. Deduction – Derivation of consequences from assumptions
   - "When it rains, the grass gets wet. It rains. Thus, the grass is wet."
   - The main type of reasoning of mathematicians
2. Induction – Creating knowledge from observed data
   - "The grass has been wet every time it has rained. Thus, when it rains, the grass gets wet."
   - Scientists are commonly associated with this style of reasoning.
3. Abduction – Creation of knowledge by imagination (also without data)
   - "When it rains, the grass gets wet. The grass is wet, it must have rained."
   - The best explanation typically used by diagnosticians (and detectives)
   - In this case the grass may have become wet for other reasons, say watered by hose

We often image that the scientific process goes on like this:

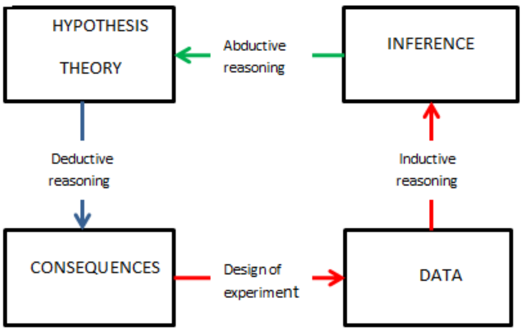

We may start with a hypothesis or a theory, possibly created by abductive reasoning without any data, e.g. just pure imagination. From the inherent assumptions we can deduce the consequences, which have to face some empirical test. We now have the background for an appropriate data acquisition, an experimental or an observational design. After having obtained the data, we obtain an inference by inductive reasoning. The result of the step may either support or be contrary to the initial hypothesis. In the latter case we may have to modify the hypothesis or the theory. The abduction step may be understood as bringing the inference to the best explanation. With an inference that is in agreement with the hypothesis or theory, it does not mean that we have proven it to be true. In fact, the logic of science literature tells us

that induction in the strict sense is impossible, expressed by Karl Popper as follows “Supporting evidence for a scientific hypothesis is merely an attempt at falsification which failed”.
When statisticians talk about statistical inference and see it as induction it defies pure logic, and uses instead a probabilistic logic. The context is typically an assumed stochastic model with some unknown features, often expressed by parameters and parameter restrictions. Data is then gathered to estimate and test the parameter restrictions.

There are several paradigms for statistical inference, the two main groupings being the classical frequentists and the Bayesians, although there are different breeds within each group. The two are likely to phrase the problem at hand differently, and also analyze the data differently. Classical statisticians typically phrase many problems in their hypothesis testing framework, seen as inductive reasoning. Most often, this is presented as testing a null-hypothesis H_0 against an alternative hypothesis H_A.
The question then arises what should be the null-hypothesis and what should be the alternative. We are typically told that the null hypothesis and its alternative should be

(i) in the specific treatment-response context
  - H_0 : No effect vs. H_A: An effect

(ii) in the general theory context
  - H_0 : Current theory vs. H_A: New theory

and the null-hypothesis is rejected and the alternative adopted when the data fall in an rejection region determined so that the probability of rejection error is reasonably low.

According to widespread philosophy of science (Popper) we can only falsify a null-hypothesis. By no rejection we have not proved that the null-hypothesis is true.

<u>Example 1:</u> Polluted river
A river is polluted and the environment authorities suspect that a nearby factory has uncontrolled releases of a toxic agents. Following the scheme above we take

H_0 : The factory is innocent H_A: The factory is responsible for releases

The burden of proof is now on the authorities, to bring forward data that allows rejection of the null-hypothesis. The type I error is then to conclude that the factory is responsible when it is innocent, and type II error is not to conclude that it is responsible when, in fact, it is. In practice the emphasis is on the null-hypothesis and controlling the type I error. The consequence is that unfair measures taken against the factory is unlikely. In a sense this is to favor the null-hypothesis. However, there may be a chance that a polluting factory goes free temporarily, due to scarce data (and high type II risk), but the verdict may be reversed with more data. This seems to be an acceptable procedure, if there is no immediate hurry of taking some action.

Contrary to this we have

<u>Example 2:</u> Polluted atmosphere

Consider the releases of an agent to the atmosphere of an activity taking place anywhere in the world, which may possibly have a long term adverse effect (e.g. fluorcarbone and ozone layer, carbone dioxide and global warming).

> H_0 : The releases of the agent have no (long term) effect
> H_A : The releases of the agent have a (long term) effect

Here we become uneasy to keep the null-hypothesis until we are fully convinced it can be rejected. This may take quite some time, and seems to violate the precautionary principle.

The classical statistical paradigm may be challenged by a number of questions:

a. Does current practice live up to its stated ideal?
b. Is the approach logically coherent?
c. Does it work in a decision context?
d. Is the role of the null-hypothesis understood?
e. Can it represent uncertainty due to lack of knowledge?
f. Can it represent subjective knowledge?
g. Can it represent knowledge accumulation?
h. Does it fit into risk management in general and precautionary thinking in particular?

Some critics answer no on all these questions, in particular the Bayesians.

The Bayesian approach is dependent on some prior belief, here illustrated in the simple case of testing two rivaling hypothesis H_1 and H_2 based on observing a variable X depending on H. Suppose that $P(x|H_i)$ is the probability of observing X=x given H_i is true (i=1,2). With prior probabilities $P(H_i)$ i=1,2, we can compute the posterior probabilities by Bayes law

$$P(H_i|x) = \frac{P(H_i)P(x|H_i)}{P(H_1)P(x|H_1)+P(H_2)P(x|H_2)} \quad \text{for} \quad i=1,2$$

This is often written in terms of the odds

$$\frac{P(H_1|x)}{P(H_2|x)} = \frac{P(H_1)}{P(H_2)} \cdot \frac{P(x|H_1)}{P(x|H_2)}$$

i.e. the posterior odds is the prior odds multiplied by the likelihood ratio for the observed data.

In some cases, the hypothesized H_1 and H_2 represent situations sufficient apart so that the difference matters, but quite often the difference is represented by a continuous parameter, so that we have to assume a prior density on the real line.

Let us contrast the Bayesian paradigm and the classical paradigm in the case of choice between a standard treatment A and a new alternative treatment B, adopted only if proven to be better. Suppose that both treatments are tried out in a setting that allows direct comparison (e.g. a randomized experiment)

<u>A classical statistician</u> will ask the question:
What is the chance to get the observed difference in favor of B, given no real difference?
<u>A Bayesian statistician</u> will ask the question:
How likely is it that B is better than A given the observed difference?

The conclusions reached by the two statisticians based on the same data may be different!

The classical statistician will presumably recommend to stick with A, unless B turns out better, and the difference cannot be attributed to chance alone (i.e. the computed P-value is low). However, it should be noticed that statistical significance is not enough, the difference should also be practically significant.

The Bayesian statistician goes more directly to the decision problem. However, in order to answer the Bayesian question we have to start with some prior beliefs about the possible differences and then use Bayes law to update this belief after observing the data, i.e. going from a prior distribution of the effect size to a posterior distribution of the effect size. The Bayesian approach has the feature that if we face additional data the former posterior now becomes the prior for the new situation.

This subjective element has long worked against the Bayesian paradigm in science, but the paradigm is now gaining ground over the classical paradigm. This is mainly due to better answers to the challenging questions addressed to the classical paradigm stated above. The three main features are:

- It can address decision problems more directly
- It is an advantage to take prior (expert) beliefs directly into account
- The addition of knowledge becomes a cumulative process

However, there are many problems with the Bayesian position too, among others:

1. How to elicit prior information?
2. How to unify differences in opinion?
3. How to represent no knowledge?
4. What if the model is likely to be wrong?

Bayesians claim to have good answers to all four questions above, but they answers are challenged, and there are differing opinions among the Bayesians themselves. For question 3, key words are non-informative priors and reference priors, and for question 4, the key word is Bayesian confirmation theory.

The common Bayesian updating feature has a possible embarrassing consequence that relates to question 4. It implies that any effect sizes imagined possible at the outset remain possible, but some may have become very unlikely. For competing models, one may also assign prior probabilities to each model, and posterior probabilities computed given the data. Contrary to classical statistics, a model is never completely ruled out by data (falsification). The possibility of rethinking the situation and replacing a model with bad fit apparently does not fit well into this scheme.

Although classical statistical theory often tries to cast hypothesis testing in a decision making context with its type I and type II error (and α and β risks), many argue that it is of little practical use. A major difficulty for many lay users of statistical theory is what to state as the null hypothesis and what to state as the alternative. Examples may be given where this is not so clear cut, and where reversing the roles may apparently give different answers, unless the whole apparatus of type I and type II errors are addressed. With these inherent difficulties in hypothesis testing, it is argued that it would be better to go for confidence intervals with a more modest ambition as decision support.

In the late 1900's the role of significance testing also became a concern for many professional scientific organizations, e.g. education, psychology, medicine, and changes in publication practice in journals were asked for. Some argue that (objective) Bayesian reporting with reference priors may resolve this issue.

Let us end this section by some thoughts about statistical methods for managing environmental health risks, where precautionary measures are typically needed. The data and modes of analysis may depend on the context, e.g.

A. Uncover exposure and its effect on human health
B. Set exposure limits
C. Monitor, predict and react

A variety of statistical methods exist, mainly within three groups

a. Regression methods
   e.g. explain mortality rates by air pollution emissions
   - yearly in different cities, controlling for different demographic characteristics
   - daily in one city with varying and occasional adverse conditions
   Estimate excess death rates (if possible)
b. Time series methods
   - Study the long term development (trend, season, randomness, episodes)
   - Study the effect of interventions
   - Monitor and perform short term predictions
c. Space-tine analysis and monitoring
   i.e. observe pollution over time at locations in a grid
   - Monitor movement (continuously) over time bases on physical models
   - Locate the sources of emissions
   - Make short term predictions

"All models are wrong, but some are useful" (George Box)

3.10 Causality and risk

Risk assessment as basis for risk treatment will typically require a causal analysis of some sort. A simple schematic tool is failure Mode and Effect Analysis (FMEA) which may involve just qualitative judgments or some simple probabilistic evaluations, for instance based on Fault Tree Analysis (FTA) or Event Tree Analysis. Another simple tool is the Cause-Effect diagram ("The Fishbone diagram"). However, this kind of analysis may be insufficient in many situations. In particular, the challenge may be high to uncover and assess cause and effect for

(i) complex systems with multiple interacting causes (and effects),
(ii) situations with only circumstantial data out of control of the observer.

Just imagine the challenges of this kind in social science and ecoscience. With this wide scope in mind In we face the more basic questions:

- What is causality?
- How can we infer causality from data?

The question of causality has been with us from the early beginnings. In the ancient world causality was not a serious problem. Gods were responsible for things to happen with a purpose, and to the extent that human beings (and animals) had a free will to cause things to happen, they were rewarded and punished by the Gods. Natural events, like storms and earthquakes, were not made causally responsible for its consequences, since they were controlled by the Gods. Even events like the outcome of the rolling of a dice were taken as a message of from God. An example from the Book of Jonah: A storm hit the ship, and in order to find out who was to blame, lots were cast! It took a long time before physical objects and processes became common as explanations for events. This mindset may have appeared first in connection with tools and equipment in everyday life. When something was broken, e.g. due to wear, it could be repaired or replaced, and it was impractical to blame it indiscriminately on God or the user. From being just carriers of credit and blame, objects now also might exhibit force, will and even purpose. Aristotle (384 BC – 322 BC) regarded purpose as the only satisfactory and complete explanation for why an object is what it is, and uncovering this was the main aim of scientific inquiry. He gave the following interpretation of causality: "All causes of things are beginnings; … that we have scientific knowledge when we know the cause; … that to know a thing's existence is to know the reason why it is". He described four classes of causes, and realized that things could be causes of one another and reciprocally causing each other, and that the same thing could be the cause of contrary effects - as its presence and absence may result in different outcomes, i.e. introduced what currently is termed a causal factor. He named two modes of causation: proper (prior) causation and incidental (chance) causation, and he spoke of both potential or actual causes and effects.

During the centuries there was not much conceptual progress, and the conceptions of causality were long pinned down by metaphysics (and still regarded as acts of God), which hampered scientific hunts for causes in the real world. It took a long time before God's role as a final cause was challenged, and gave room for human knowledge. A breakthrough came with Galileo (1564-1642). In his work Discorsi from 1638 he states two maxims:

- One, description first, explanation second – that is , the "how" recedes the "why", and
- Two, description is carried out in the language of mathematics; namely equations

According to Galilei we should not ask whether an object is falling because it is pulled from below or pushed from above, but rather ask how well we can predict the time it takes for the object to travel a given distance, and how that varies under varying circumstances. His general ideas became rapidly widespread in the natural sciences, and science went from speculative to empirical, which led to the many discoveries of physical laws in the centuries to come. However, most philosophers were reluctant to give up striving for finding causal explanations behind the successful empirical discoveries and mathematical descriptions.

The philosopher David Hume (1711-1776) went one step further from Galilei by removing the "why" as in Galilei's first maxim as an objective. He advocated that any connection between cause and effect is not due to reason, and that effects are distinct events from their causes, since we can always imagine one such event occurring and the other not. So, according to Hume, causes and effects are discovered, not by reason but through experience, when we find that particular objects or events constantly follow or are adjacent to one another. Hume went one step further, by disclaiming any attempt to explain the link between causes and effects in terms of powers, active forces, like the power of God to cause things to happen. Thus he secularized completely the notion of causality. This was revolutionary thoughts at the time and, for instance, Kant (1724-1804) argued that people already possessed innate assumptions related to causes.

By his negative initial argument, Hume seemingly had no link between the past and the future. He nevertheless recognized that the concept of causality involved the idea of necessary connections. Where should this idea come from, if there is no perception of necessity in causal sequences? Instead of taking the necessity to be a feature of the natural world, Hume took it as a feature arising within the human mind, influenced by the observation of regularity in nature which formed an expectation of the effect, when the cause is present. The "essence of necessity" is, according to Hume, "something that exists in the mind, not in the objects". His lesson was that a priori reasoning and arguments leads us astray: "It is only experience which teaches us the nature and bounds of cause and effect, and enables us to infer the existence of one object from that of another". Since we all have limited experience, our conclusions should always be tentative, modest, reserved, cautious (a position he named *mitigated skepticism*).

Note: Efforts to reconcile the views of Hume and Kant exist, recently within psychology by Patricia Cheng (1997). According to her power PC theory, people filter their observations of events through a basic belief that causes have the power to generate (or prevent) their effects, thereby inferring specific cause-effect relations. The theory assumes probabilistic causation, and may be linked to other developments in this area, among others within a Bayesian framework.

The philosopher George Berkeley (1685-1753) had clearly understood that correlation does not imply causation, but we had to wait for some pioneer statisticians to get a better grip on this.

Francis Galton (1822-1911) and Karl Pearson (1857-1936) exposed the limitations of statistics in inferring causality from data and advocated the concept of covariation, in terms of contingencies and correlation. Then causation was given no place in statistical theory. Later Ronald A. Fisher (1890-1962) developed the theory of controlled experiment, which provided an opportunity to reveal causal effects, but for many decades causality was given no role beyond that.

Causality played little role in the natural sciences as well. Physical formulas, like Newton's second law F=m/a, have no direction in themselves. Although scientists may talk in causal terms at coffee tables, the scientific reports are not phrased in such terms. An effort to bring causality into the natural sciences was made by the famous physicist Max Born (1880-1970), who characterized causality as follows:

i. <u>Causality</u> postulates that there are laws by which the occurrence of an entity B of a certain class depends on the occurrence of an entity A of another class, where the word entity means any physical object, phenomenon, situation, or event. A is called the cause, B the effect.
ii. <u>Antecedence</u> postulates that the cause must be prior to, or at least simultaneous with, the effect.
iii. <u>Contiguity</u> postulates that cause and effect must be in spatial contact or connected by a chain of intermediate things in contact.

New insights to relativity and quantum mechanics have forced physicists to abandon these postulates as description for what happen at the most fundamental level, but they may still be regarded valid at the level of human experience.

Strangely enough, we had to wait to the last decades of the 20th century to get a better grip on probabilistic causation and causal inference. For a long time a widely held view was that statistics alone cannot reveal causation, but just correlation. In recent years this has changed, partly due to new efforts to define concepts of statistical causality useful for observational studies, like Granger-causality in a time series context (named after the econometrician Clive Granger 1934-). This is an effort to entangle causation from correlation, based on prediction of time series from other series. Worth mentioning is also the concept of local independence in Markov models (Tore Schweder). Among areas that inspired this developments are epidemiology, social science and environmental science.

The developments in graph-modeling with computerized implementations, e.g. like Bayesian belief networks (see next section), has widened the opportunities for causal analysis and inference. This has been applied in many areas, among others as decision support in medicine.

More recently, a group of ecoscientists have suggested methods for distinguishing cause-and-effect from misleading correlation (George Siguhara et.al, 2012). Their method is named "convergent cross mapping", CCM for short. The CCM method apparently has demonstrated some success, as in the following example. However, the method has already received some critics (as most methods do), and its potential remains to be seen.

<u>Example</u>
The population sizes of anchovies and sardines are known to move in opposite directions, i.e. there is a negative correlation, but this correlation is spurious, i.e. no causation. However the

CCM method was able to demonstrate that see temperature was a hidden driving force for size changes of both fish stocks, despite the fact that temperature was not correlated with any of the two fish populations.

With the challenge of this example in mind we go on to explore some basic ideas related to causation, and mainly to probabilistic causation, which seem inevitable.

Types of causes

Causes are often distinguished into two major types:

- **Necessary cause:**
  If *A* is a necessary cause of B, then if *B* has occurred then necessarily A also has occurred. However, the presence of *A* does not necessarily mean that *B* will occur.

- **Sufficient causes:**
  If *A* is a sufficient cause of *B*, then the presence of *A* necessarily implies the presence of *B*. However, the presence of *B* does not imply the presence of *A*, since another cause *A'* may alternatively cause *B.*

We also need to handle the possibility the outcome is not due to a single cause, and talk of

- **Contributory causes:**
  A is contributory to *B* when *A* precedes *B*, and altering *A* alters the effect *B*. Here we allow to interpret "altering" in a probabilistic sense.

A contributory cause may be neither necessary nor sufficient. Example: In the context of epidemiology, it does not require that all individuals which possess the contributory cause *A* experience the effect *B*, and it does not require that all individuals who are free of the contributory cause *A* will be free of the effect *B*.

Example: Short circuit as cause for house burning down (J. L. Mackie).
Consider the collection of events: the short circuit, the proximity of flammable material, and the absence of firefighters. Together these are sufficient but not necessary for the house to burn down (since many other collections of events could have given the same outcome). Within this collection, the short circuit is an insufficient but non-redundant part, since the short circuit by itself would not have caused the fire, but the fire would not have happened without it, everything else being equal. This special kind of contributing cause satisfies the so-called INUS condition: Insufficient and Non-redundant parts of Unnecessary, but Sufficient causes (often revealed by a fault tree).

Deterministic vs probabilistic causation

Causality has in practice (and in logic and philosophy) often been interpreted in a deterministic context, meaning that if *A* causes *B*, then *A* must *always* be followed by *B*. In this sense, war does not cause deaths, nor does smoking cause cancer. This does not seem to be fruitful. Consequently, we may turn to the notion of probabilistic causation. Informally,

A causes *B* in the probabilistic sense if *A*'s occurrence increases the probability of *B*.

In some cases this may reflect imperfect knowledge of the system under study, whether it is regarded as deterministic or not, a knowledge not worth chasing for. In other cases it may reflect a genuine stochastic system (leaving out the philosophical and theological issue whether genuine randomness exists).

In order to analyze causation in science, business and everyday life, probabilistic causation seems inevitable. This leads to many challenging issues that need clarification, mostly provided by workers in fields like statistics and information science (and psychology), or within specific application fields itself, like physics, engineering, biology and medicine. Workers in specific fields may regard it useless to adhere to any universal philosophical conception of causality, whether probabilistic or not.

In science the mode of observation is regarded decisive for concluding causality. In general we have to differentiate between

- Controlled experiments,
  - the observer can control and vary the causes keeping "the environment" fixed
  - e.g. natural sciences, medicine, agriculture
- Quasi-experiments
  - the observer can vary some causes, but "the environment" is not quite fixed
  - e.g. social experiments (before and after action)
- Observational studies,
  - the observer have no control over the causes
  - e.g. epidemiology

In a well designed controlled experiment there is good opportunity to infer causal effects, but less so in an observational study. Here we may observe co-variation that is not linked to the causal issue to be explored, but to something else, e.g. self-selection. In the case of smoking and lung cancer there is a difference between

A. P(Cancer|Smoking)
B. P(Cancer|Forced to smoke)

The first can be estimated directly from an observational study, but does not represent the causal effect of smoking, since we cannot a priori disregard the possibility that people prone to cancer are more likely to pick up the habit of smoking. In a controlled experiment we could in principle assign people randomly to each group, smoker and non-smoker, and imagine to estimate the probabilities for both groups, and compare the difference. In practice such a controlled experiment was not possible (time limit, cannot force) and science had to relate to observational studies (and to controlled experiments on rats exposed/not exposed to nicotine). The causal conclusions from these observational studies were disputed, but eventually generally accepted. However, the debate on these issues is still ongoing.

This led to more basic research on what can be causally inferred from observational studies, and also to development of more practical analytic tools to help uncover and quantify causal effects, mainly developed within the artificial intelligence (AI) community. Among such tools are

- Causal calculus
- Structure learning

Both involve a graphical structure of possible causal relations among variables. Causal calculus is usually based on so-called Bayesian Networks, a graph structure with nodes and directed arcs representing assumed causal relationships, where conditional probabilities are specified. From this one can compute interventional probabilities in a consistent manner (e.g. a removed cause). The method allows unobservable variables. Structure learning may depart from a so-called skeleton, a graph without causal arcs between the variables, and then observable statistics may, under some assumptions, help to uncover the directions of some of the arcs. Other types of structural learning are based on search among many possible causal structures, and remove those who are strongly incompatible with the observable correlations. Such approaches are disputed as being unscientific, but have found practical application in many areas.

In practice we often hear statements that are easily taken as indicative of a cause relationship, but may be just founded on an observed association or just being a conditional statement "If...then". This may lead to confusion, in particular if probabilities are involved.

The philosopher David K. Lewis (1941-2001) suggested that all statements about causality can be understood counterfactual. Example: "John's smoking caused his premature death" is equivalent to saying "Had John not smoked he would not have prematurely died". The computer scientist Judea Pearl (2000) has demonstrated that the translation of causal statements into counterfactual statements is both feasible and operationally useful. For instance, one can in principle compute the probability that John would be alive had he not smoked given that, in reality, John did smoke and did die.

Note: Suppose that John did smoke and died as a result of his smoking. However, he had an enemy who would have killed him shortly after anyway. The counterfactual position now faces a theoretical problem in order to claim that smoking caused John's death since, had John not smoked, he still would have a premature death.
(extensively philosophical discussions of this example exist).

We will here give a brief impression of ideas related to probabilistic causation:

Probabilistic causation

Let *A* and *B* represent events or factors that are potentially causally related, both with probabilities strictly between zero and one. Suppose we say that A causes B if and only if the probability of B is higher if we know that A happened than when we know that A did not happen, i.e.

$$A \text{ causes } B \leftrightarrow \mathrm{P}(B \mid A) > \mathrm{P}(B \mid \text{not-}A).$$

Note: Here we have in mind promoting cause, if the inequality sign goes the other way, we may talk about inhibiting cause.

It is easy to check that this is equivalent to

$$\mathrm{P}(B \mid A) > \mathrm{P}(B) \text{ and also to } \mathrm{P}(A \,\&\, B) > \mathrm{P}(A)\cdot\mathrm{P}(B).$$

We see that the latter expression is symmetric in A and B, and we therefore also have that $\mathrm{P}(A \mid B) > \mathrm{P}(A \mid \text{not-}B)$, so we just as well could have said that "B causes A". All these relations therefore express nothing more than an association. In our case, promoting cause, we have positive association. In the case of inhibiting cause, with the inequalities going the other way, we have negative association, and we may or may not have causality. In the case of all relations being equalities, we have that A and B are independent, most often expressed symmetrically by

$P(A \& B) = P(A) \cdot P(B)$. Are we then justified to say that A and B are not causally related? Not at all! They may in fact, still be causally related.

The above may happen in the following context: Consider the case that *A* and *B* are both caused by a third factor *C*, as indicated in the left diagram. It turns out that we may have P(*B* | *A*) > P(*B* | not-*A*) even though *A* does not cause *B.*

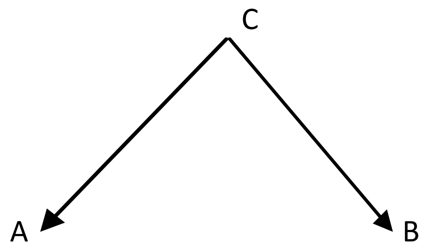

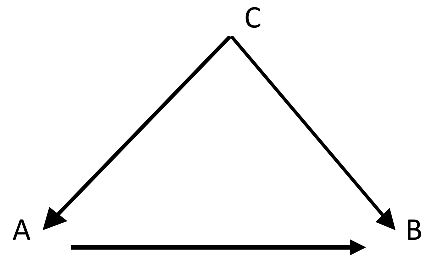

A simple example may serve to illuminate this*:*

Let A be that an individual has yellow-stained fingers, and B that the individual has lung cancer. Then we expect that P(*B* | *A*) > P(*B* | not-*A*), the reason being that those with yellow-stained fingers are more likely to suffer from lung cancer than those without, is that smoking (C) tends to produce both effects. Because individuals with yellow-stained fingers are more likely to be smokers, they are also more likely to suffer from lung cancer. This is called *spurious correlation*.

In the right diagram we illustrate that C both causes B directly and causes B via A. Using the notion above we should therefore have we P(B|C) > P(B|not-C) and P(B|A&C) > P(B|A&not-C), Nevertheless we may still have that P(*A* & *B*) = P(*A*) · P(*B*), so that A and B is independent. Observed non association does not therefore exclude the possibility of a causal relationship in the sense above.

Consequently, the above definition of causation is not workable without further qualifications. For situations where the events A and B can be attached to a time scale, so that A happens before B, we have the following possibility:

<u>Definition:</u> A before B and causes B ↔ (is equivalent to saying)

i. P(B | A) > P(B | not-A)
ii. No C exists, before or simultaneously with A, with P(*B* | *A* & *C*) = P(*B* | *C*),

For the case that P(*B* | *A* & *C*) = P(*B* | *C*), *C* is said to screen *A* off from B, and condition ii is named the 'No Screening Off' condition (Hans Reichenbach).

When P(*A* & *C*) > 0, screening off is equivalent to P(*A* & *B* | *C*) = P(*A* | *C*)·P(*B* | *C*), which means that A and B are independent given C. This means that *C* renders *A* probabilistically irrelevant to *B (*in fact to each other).

Example: Suppose that smoking (*C*) causes both yellow-stained fingers (*A*) and lung cancer (*B*). Then smoking will screen yellow-stained fingers off from lung cancer: Given that an individual smokes, his yellow-stained fingers have no impact upon his probability of developing lung cancer.

The second condition in the definition was added to eliminate cases where spurious correlations give rise to factors that raise the probability of other factors without causing them. But will the added condition completely resolve the problem of spurious correlations? No! Spurious correlations can also give rise to cases where a genuine cause does not raise the probability of its effect at all, i.e. a cause need not satisfy the first condition.

Example: Suppose that C denotes physical exercise, and that (perhaps contrary to reality) that smokers are much more likely to exercise as well. Smoking (A) is a cause of heart disease (B), but suppose that exercise is an even stronger preventative of heart disease. Then it may be that smokers are, over all, less likely to suffer from heart disease than non-smokers, i.e. $P(B \mid A) < P(B \mid \text{not-}A)$. On the other hand, if we condition on the fact that the individual exercises, the inequality is reversed: $P(B \mid A \,\&\, C) > P(B \mid \text{not-A} \,\&\, C)$, and if we condition on the fact that the individual does not exercise, the inequality is reversed as well: $P(B \mid A \,\&\, \text{not-}C) > P(B \mid \text{not-}A \,\&\, \text{not-}C)$. Such reversals of probabilistic inequalities may occur in many contexts, and may present itself as a paradox, often named "Simpson's Paradox."

By doing controlled experiments and keep confounding factors fixed, we can avoid the trap of spurious correlations. In the case of observational studies, we may imagine that we condition on confounding factors, we keep them fixed. Investigating whether A causes B, keeping a specified set of other factors fixed, will be named a test situation T for A and B.

Definition

A causes *B* $\leftrightarrow$ $P(B \mid A \,\&\, T) > P(B \mid \text{not-}A \,\&\, T)$ for every test situation T for A and B.

i.e. causes must raise the probability of their effect in test situations.

A question here is which factors are necessary and sufficient to involve in test situations. We saw that the true causal impact of smoking for lung cancer was revealed when we held exercise fixed, either positively (conditioning on C) or negatively (conditioning on not-C). This suggests that in evaluating the causal relevance of A for B, we need to hold fixed other causes of B, either positively or negatively. Should every such cause be included? No!

Example (cont'd): Consider smoking (A) and lung cancer (B) and a causal intermediary (C), say the presence of tar in the lungs. If A causes B exclusively via C, then C will screen A off from B: given the presence (absence) of carcinogens in the lungs, the probability of lung cancer is not affected by whether those carcinogens got there by smoking (are absent despite smoking).

Thus we will not want to keep fixed any causes of B that are themselves caused by A.

We have talked about promoting cause and inhibiting cause (preventers). Now we understand that in order to give this label, we have to be sure that the effect goes in the same direction in all test situations. This may not be the case. We may have causes that affect the probability of its 'effect' in different ways in different test situations. Such causes may be named an 'interacting' or a 'mixed cause'. It should be apparent that when constructing test situations for A and B one should also keep fixed preventers and mixed causes of B that are independent of A.

Above we have discussed causality in term of events and probabilities, i.e. in a binary context, where the events A and B occur or not. We may want to extend the ideas to causal relationships between non-binary variables X and Y, such as caloric intake and blood pressure. In assessing the causal relevance of X for Y, we will need to hold fixed the values of variables that are independently causally relevant to Y. In principle, there are infinitely many ways, in which one variable might depend probabilistically on another, and infinitely many ways of keeping a specific test situation fixed. The basic features of the theory generalize to non-binary variables, but we do not get the neat classification of causal factors into promoting and inhibiting causes.

Admittedly, there are some problems with linking causality too strong to precedence in time. It may not be easy to decide whether a property or type of event is preceding another, since in many cases both may develop concurrently. This may invite the researcher to postulate causal relationship based on a faulty perception of the time scale.

A comprehensive and authoritative book on causation is authored by Judea Perl (2000).

When we assign probabilities to several interrelated events in risk situations the probabilities should be consistent, i.e. throughout fulfill the common rules of probability. This is hard in situations with possible spurious correlation, and for risk models in practice we typically have to go from some agreed causal scheme. Then Bayesian net ideas may come to help, which is the theme of the next section.

Science usually tries to explain an outcome state by a single or a combination of few interacting causal factors, and consequently adopt research methods suited to uncover the one or the few factors. Then it may be possible to interfere in the causal chains. Linear models are often sufficient. This reductionist view has been successful in many areas, and statistical theory may provide modelling opportunities and methods to avoid confounding factors.
Within some areas, like health and environmental science, this may not always be feasible. There may be no single causal factor, but a multitude of factors hard to order in causal chains. A health risk factor can be introduced to the human body at an early age, stay and develop concealed for years, in conjunction with other factors, and possibly manifest itself as an illness later in life.

<u>Example</u>. Tobacco smoke contains about 4000 known chemical substances, and more than 100 of them are identified as toxic. The process leading to cancer to some smokers and not to others is still not understood.

In such areas, a more holistic approach may be needed. This may imply that a details contextual description of a few carefully selected cases is favored to repeated measurement of few expected key factors that relates to some measured outcome. However, the challenges are many, and be prepared to face irreducible uncertainties, non-linear dynamics and emerging properties. Theory and methods that deal with this are still not well developed, and it is tempting to fall back to classical methods. In some sciences, there are disputes between the reductionist/variable view and the holistic/case view. This is so in environmental science.

3.11 Bayesian belief methods

Bayesian classifier

A frequently occurring situation is to decide, based on observed data whether an activity, has had an impact or not. The context may be within the workplace, in questions of public health or for an environment issue. Let us as illustration consider the latter contest.
Think of a possible contaminated site after a spill nearby, and with two states (impact, no impact). In many cases there may be some impact, but not of practical significance, so that the conclusion may be one of (impaired, not impaired) which may lead to a choice between two decisions (remediate, do not remediate). Our discussion with just two states may then apply for such cases as well. It may be convenient to think in terms of odds:

$$O(impact) = \frac{P(impact)}{P(no\ impact)}$$

If the probability of impact is greater than the probability of no impact then O(impact)>1.
With no data, the probabilities have to be based on prior information or beliefs, and will thus to a large degree subjective and personal.

We want to be able to evaluate the odds given the data. Often data may come from various sources, and we want to combine them in a consistent manner. It may not be easy to describe impact, and a way out is to observe the same quantities at some reference locations, known to be unaffected by the spill.

From Bayes' law we get:

$$\frac{P(impact \mid data)}{P(no\ impact \mid data)} = \frac{P(impact)}{P(no\ impact)} \cdot \frac{P(data \mid impact)}{P(data \mid no\ impact)}$$

We therefore have

$$O(impact \mid data) = O(impact) \cdot LR(data)$$

Where LR is the likelihood-ratio of the data (also named Bayes factor) given by

$$LR(data) = \frac{P(data \mid impact)}{P(data \mid no\ impact)}$$

Note that if we a priori believe that the stated impact and no impact are equally likely the O(impact)=1, and so O(impact|data)=LR(data).

Even with data, we do not have the probability of impact itself, only the probabilities derived from a model for how the data may look like at impacted sites and not impacted sites.

For environmental problems, there may not be simple approaches for estimating these quantities. Building a model for the impacted or unimpacted sites requires information on level and variation (distribution) and factors that might influence the observations. At the modeling stage the Bayes factor may involve unknown parameters, for which we may attach an uncertainty distribution which may be integrated out.

For those who prefer to state the strength of evidence for impact in verbal terms, we offer the following (due to Good, 1988):

Strength of evidence	Weak	Moderate	Moderate to Strong	Strong
Bayes factor LR	<5	5-10	10-100	>100

Now consider that the data is from two, possibly correlated, sources, we write $X=(X_1, X_2)$. Then the odds for A=impact may be written as

$$O(A \mid X_1, X_2) = O(A) \cdot LR(X_1) \cdot LR_1(X_2)$$

where the second LR-factor is the likelihood-ratio for observing X_2 given that X_1 is observed. If the two sources are independent, this is the ordinary LR-factor. In case of dependency, the second factor should be diminished, and one (somewhat ad hoc) suggestion is to compute

$$O(A \mid X_1, X_2) = O(A) \cdot LR(X_1) \cdot LR(X_2)^w$$

where w is a weight between 0 and 1 that accounts for positive correlation. The weight may be derived from different principles. In the case of two numerical variables we may use $w=1-R^2$, where R^2 is the coefficient of determination by regressing X_2 on X_1. This extends to more than two lines data:

$$O(A \mid X_1, X_2, \dots, X_r) = O(A) \cdot \prod_{i=1}^{r} LR(X_i)^{w_i}$$

with $w_1=1$ and $w_i=1-R_i^2$ where $R_i^2=R^2(X_i \mid X_1, X_2, \dots, X_{i-1})$ for i>1 is the coefficient of determination by regressing sequentially for X_i on $X_1, X_2, \dots, X_{i-1}$ for i>1. This introduces some arbitrariness by the order we take into account the different data.

<u>Remark.</u> By taking the logarithm of the odds we get a linear expression in the Bayes factors. It then seems natural to use log(LR) as a weight of evidence measure, which then becomes additive

Bayesian Belief Networks

For complex systems there are typically many circumstances and events that may lead to hazards. The uncertainties of these events may be of varied nature, and in order to describe our uncertainties about them, several methods are available: Probabilities, Three-value logic (true, false, perhaps), Fuzzy set logic, Possibility logic, non- monotone reasoning. The challenge is to keep track of many interrelated events, and provide a suitable terminology and methodology for analysis. One possibility made popular by available software during the last decade is Bayesian Belief Networks (in short Bayes Nets), which is based on probabilities.

Bayes nets may answer questions like:

- What are the chances for this problem to arise, given these symptoms?
- Is it something peculiar with this observation?
- What should we do further when we have observed this?

The problem is to specify joint probabilities in a systematic and consistent manner, in accordance with the basic properties and rules of probability calculus. Here Bayes nets come to help by representing causes and effects via an intuitive graphical representation, and accommodates situations where some information is known and some is uncertain vague, incomplete, conflicting or unavailable.

Each variable in a Bayes net is represented by *nodes,* and each node has a set of probable values for each variable, named *states*. The nodes are connected by edges with an arrow to indicate the direction of influence. Relationships between some of the states are expressed by conditional probability relationships depending on the stated causalities. Nodes can represent any kind of variable, observable or not, and even a hypothesis. Efficient computer algorithms are available that perform inference and learning in Bayesian networks. Bayes nets have found wide applications in many fields: Medicine, finance, oil exploration, ecosystem modeling. The main tasks are merging of expert opinions, automatic diagnosis, monitoring and alerting. We will illustrate some of the features of a Bayesian Net by an example with nodes representing events with just two outcomes.

Example Bayesian Net

Al and Ben both go to work in the morning by a 7.15 AM bus on different bus lines to reach their jobs on time at 8.00 AM. They may both be late due to slow traffic, in the winter usually snowfall, that typically may affect both bus lines. Sometimes Ben is late due to a special cause: He is slow in the morning and did not catch the bus, and have to take the next bus at 7.30 AM. Al however is an early riser, and never misses his bus. The causal picture for being late may look like this:

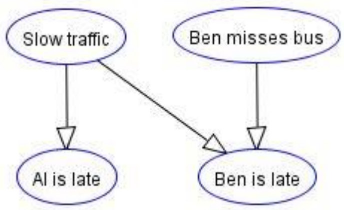

Here each box corresponds to an event with outcome “True” or “False”. We could alternatively represent the outcomes by an event tree, but this becomes unnecessary involved.
Suppose we have the following probabilities for Al being late due to traffic jam:

Slow traffic	Al late	
	True	False
True	0.5	0.5
False	0.1	0.9

and the following probabilities for Ben being late due to him missing bus and/or slow traffic:

Ben misses bus	Slow traffic	Ben late	
		True	False
True	True	0.8	0.2
True	False	0.4	0.6
False	True	0.5	0.5
False	False	0.1	0.9

Suppose then that the (prior) probabilities of “Slow traffic” and “Ben misses bus” are respectively:

Slow traffic	Priors
True	0.2
False	0.8

Ben misses bus	Priors
True	0.1
False	0.9

The fact that there is no connection or no common parent node to “Slow traffic” and “Ben misses bus” means that these two events are taken to be independent. . From all of this we may derive various probabilities by means of common probability calculus rules. The probability of “Al is late” and the probability of “Ben is late” may be found by conditioning. Various conditional probabilities, like the probability that “Al is late” given that “Ben is late” may be found using Bayes´ law. Computations like this are trivial but tedious, and with more nodes and more involved structure, it is easy to loose track of what to do. Fortunately, we are helped out by user-friendly software. Here is an example of output from a freeware named “Bayes”:

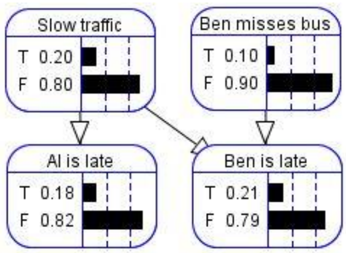

We see that the probabilities of “Al is late” and “Ben is late” are 18% and 21% respectively. If we get to know that Al is late, i.e. observed value is True (T), we get conditional probabilities as follows:

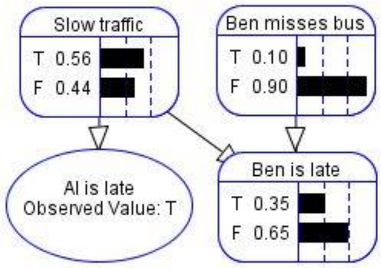

We see that, given “Al is late”, the probability of “Slow traffic” is modified to 56%, the probability of “Ben is late” to 35%, while the probability of “Ben misses bus” is unchanged. On the other hand, if we get to know that “Ben is late”, i.e. observed value is True (T), we get conditional probabilities as follows:

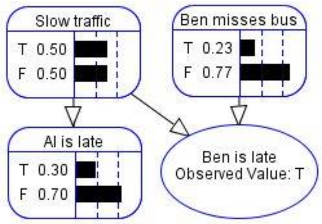

We see that, given “Ben is late”, the probability that he missed the bus is 22%, the probability of “Slow traffic” is 50%, while the probability of “Al is late” now becomes 30%.

We have tacitly assumed that the slow traffic probabilities refer to the situation for the 7.15 AM bus and relevant for both lines. However, it may be that the risk of slow traffic increases towards 7.30 AM. In case of Ben missing the bus, one should therefore increase his late probabilities due to “Slow traffic” at 7.15 AM. Alternatively, one should redesign the causal graph with separate nodes for the two bus departures.

Exercise

a. Derive some of the probabilities above by hand.
b. Modify the graph to account for differences in the traffic on the two lines and departures.
c. How do we express the situation where they come to work with the same bus?

Combination of expert judgments: Bayesian

Let X be an unknown quantity and $x_1, x_2, \dots, x_n$ be estimates of X given by n experts in an expert panel. Suppose we know the probability distribution of the expert judgments given the true value X=x

$$p(x_1, x_2 ,\dots, x_n \mid x)$$

If the decision maker has a prior distribution p(x) for X, the posterior distribution for X given $x_1, x_2 ,\dots, x_n$ is given by the Bayes formulas as

$$p(x \mid x_1, x_2 ,\dots, x_n) \quad \text{prop.} \quad p(x)\, p(x_1, x_2 ,\dots, x_n \mid x)$$

where prop. means proportional to. If the experts judge independent of each other we have

$$p(x_1, x_2, \ldots, x_n \mid x) = \prod p(x_i \mid x)$$

where the right hand side is the product of the probability distributions of each expert.

To get further we have to make assumptions of each of the factors. A possibility is to adopt an additive model, where we write

$$x_i = x + e_i$$

i.e. the estimate of expert no.i is the true unknown value x plus an error e_i.

If we assume that this error is random and normally distributes with expectation m_i and standard deviation s_i . Here an m_i different from zero may represent potential systematic bias of expert no.i, and standard deviation s_i how precise the expert is beyond the bias.
If we at the same time assume that the prior distribution is normal with expectation x_0 and standard deviation s_0 , it follows that the posteriori distribution is also normal with expectation

$$E(X \mid x_1, x_2, \ldots, x_n) = \sum w_i (x_i - m_i) \qquad (m_0=0)$$

where ∑ means the sum over all 0,1,2, .., n and $w_i = s_i^{-2} / \sum s_k^{-2}$ for i=0,1,2,…,n. A similar formula may be given for the posteriori standard deviation, which will express how uncertain we are after the expert judgment.

The method requires a prior probability distribution from the decision maker. Even if this challenge with the assumption of normality is reduced to the specification of a prior expectation and standard deviation, this is something that many decision makers would not like to do, maybe by saying ”I have no idea!” In this case a possibility is to use a so-called non-informative prior distribution.

Bayesian principles for the combination of expert judgments may be applicable in many contexts, in the judgment of measurable quantities, as well as in the judgment of probabilities. If the judgments are part of a larger context as one of many, for example as part of a decision analysis the challenge is bigger.

4 Cases

(supplied separately)

5 Literature

Aven: Foundations of Risk Analysis: A knowledge and decision-oriented perspective, Wiley, 2003.

Aven: Risk Analysis: Assessing Uncertainties Beyond Expected Values and Probabilities, Wiley, 2008.

Aven: Quantitative Risk Assessment: The Scientific Platform, Cambridge University Press, 2011.

Aven & Vinnem: Risk Management with Applications from the Offshore Petroleum Industry, Springer 2007.

Bedford & Cooke: Probabilistic Risk Analysis: Foundations and Methods. Cambridge University Press, 2001.

Klinke & Renn: A New Approach to Risk Evaluation and Management: Risk-Based, Precaution-Based, and Discourse-Based Strategies, Risk Analysis, vol. 22 No.6, 2002.

Martuzza & Tickner (Eds): The precautionary principle: protecting public health, the environment and the future of our children, WHO Europe 2004.

NENT: The Precautionary Principle: Between Research and Politics, Report from The National Research Ethical Committee for Natural Science and Technology (NENT), National Committee for Research Ethics in Norway 1997.

NENT: Risk and Uncertainty – as a Research Ethics Challenge, Report from The National Research Ethical Committee for Natural Science and Technology (NENT). National Committee for Research Ethics in Norway, 2009.

Pearl: Causality-Models, Reasoning and Inference, Cambridge University Press 2000.

Perrow: Normal Accidents: Living with High Risk Technologies, Princeton University Press 1984

Perrow: The Next Catastrophe: Reducing Our Vulnerabilities to Natural, Industrial, and Terrorist Disasters. Princeton University Press 2007.

Renn: Risk Governance: Towards on Integrative Approach, White paper no. 1, International Risk Governance Council 2005.

Sklet: Methods for accident investigation, Reliability, Safety and Security Studies at NTNU, 2002

UN/ISDR: Living with Risk, A global review of disaster reduction initiatives, 2004 version.
The UN Inter-Agency Secretariat of the International Strategy for Disaster Reduction.

Vose: Risk Analysis: A Quantitative Guide, Wiley 2000.