

Society for Risk Analysis

Risk Analysis Foundations

May 7, 2015

Preface

The idea of this paper is to prepare a document which reflects on key scientific pillars of risk analysis, the core of our scientific, regulatory and technical field, the elements that unify our professional discipline, with both current and future perspectives. The document does not aim to provide some agreed basic pillars for risk analysis as a field. Rather, it wants to serve as a stimulator of discussion about what these pillars can and should be. The goal is, then, not to establish a unified consensual perspective on the area of risk analysis– providing the right views on specific definitions and methods – but to highlight what are some of the main issues, their nature and rationale, and lay the ground for discussing them. We do not intend this to convey an “official” SRA opinion on the topics addressed; rather, this document reflects the thoughts of the authors. Indeed, not all of the authors agree with everything in this paper, but we all support the publication of this paper as a means to foster reflections and exchange of ideas, hoping that this will help to advance the foundations of the field of risk analysis.

The discussion is open to the contribution of all members of any Specialty Group, and others. We believe that many can benefit from reading the paper, a rather short summary of some open-minded reflections on fundamental concepts and principles of our risk analysis field.

At this stage the paper includes discussions of the following topics:

- 1) Risk analysis and science
- 2) The risk concept
- 3) Risk management principles
- 4) Uncertainty in risk analysis
- 5) Confronting deep uncertainties, surprises and the unforeseen
- 6) Reliability, validity and trustworthiness of risk analysis methods and results (including suggestions of how to make them more trustworthy)
- 7) The future of risk analysis: meeting the challenges. Emerging trends.

We aim at including other issues at a later stage. The document is planned to be updated from time to time to reflect the ongoing discussion, addressing comments and suggestions made. Please contact terje.aven@uis.no if you have some ideas and/or views about the document and potential future work.

The paper has been a result of an initiative taken by the Specialty Group on Foundational Issues in Risk Analysis (FRASG) of the Society for Risk Analysis (SRA). It has been developed by the following group of experienced and active researchers:

Terje Aven (leader)
Henning Boje Andersen
Tony Cox
Enrique López Droguett
Michael Greenberg
Seth Guikema
Wolfgang Kröger
Ortwin Renn
Enrico Zio

The work has been carried out during the first part of 2015 and will be presented for the SRA Council in June 2015.

1 Risk analysis and science

The risk analysis field is represented in many scientific journals and conferences, and there are quite a few professorships and university programs covering risk assessment, risk management and related areas. Thus, it can be argued that the risk field is a science since it has the standard characteristics of a science. In fact, the field is based on the same type of universal criteria as other scientific fields/disciplines, for the judgment of the quality of its scientific work.

This scientific field/discipline of risk analysis develops:

- i. knowledge about risk-related phenomena and processes, triggering events and event chains, etc. (for example the consequences of using a specific drug in a medical context or the damages to the environment caused by an oil spill in a specific type of coastal area)
- ii. new (modified) concepts, theories, frameworks, approaches, principles, methods and models to understand, assess, characterize, communicate and (in a wide sense) manage/govern risk, in general and for specific applications (the instrument part). The former type of knowledge is usually obtained by combining insights from different disciplines, for example medicine, biology and ecology, with various formal modeling approaches, most commonly traditional statistical analysis involving hypothesis testing or Bayesian analysis. The latter type may, for example, cover a new way of describing uncertainty in risk assessment.
- iii. data bases and methods to evaluate these data

Science is a means to produce knowledge – the most reliable statements and justified beliefs that can be made at the time being – on subject matters covered by the community of the field. In simple words, we can say that the knowledge part of this field is about understanding the world (in relation to risk) and *how* we can better understand, assess and manage this world (in relation to risk).

The distinction between (i) and (ii) types of knowledge leads us into the common distinction between phenomenological and descriptive studies (including empirical research: knowledge gained by observations or experiments) on the one hand and normative theory based on conceptual tools on the other. One may ask whether this normative part is about gaining knowledge, as it is more about providing arguments for what to do, for example how we should understand the risk term and best analyse risk. When entering the normative dimension, the science and knowledge dimension may seem lost, at first glance, but this is not the case. Provided that the analysis is solid/sound (follow standard protocols for scientific work like being in compliance with all rules and assumptions made, the basis for all choices are made clear, etc), the risk field gains new insights by seeing the arguments leading up to the recommendations and conclusions. There are explicit or implicit value judgments in the process, yet the process and argumentation could enrich the existing insights and literature. For the work, it is essential to acknowledge the value judgment part – that the analysis provides the authors' conclusion based on a set of arguments; there could be other weights given to the various arguments, and hence other conclusions. This is, however, not only

relevant for this type of contributions: all scientific work is to some extent framed and influenced by the authors' choice of concepts, methods and models, and has to be interpreted with this in mind.

To what extent risk assessments and other risk studies and works are scientific is another question. Some of the earliest contributions to the discussion go back to Alvin M. Weinberg and Robert B. Cumming, and the editorials of the first issue of the *Risk Analysis* journal in relation to the establishment of the Society for Risk Analysis. These authors concluded that risk assessment is not a scientific method per se, when the reference is the "traditional scientific method", standing on the pillars of accurate estimations and predictions. They recognized and asserted the fact that there are, and will always be, strong trans-scientific elements in risk assessment (questions which can be asked of science but yet cannot be answered by science, for example predictions of rare events where the uncertainties are very large). However, a risk assessment can also be seen as a tool used to represent and describe knowledge and lack of knowledge, and then other criteria need to be used to evaluate whether the assessment is a scientific method. Adopting a broad interpretation of scientific methods, risk studies are scientific if the work is solid/sound (see above), is relevant and useful (support decision making), and meet the criteria of reliability and validity. The reliability requirement relates to the extent to which the risk assessment yields the same results when repeating the analysis, and the validity requirement refers to the degree to which the risk assessment describes the specific concepts that one is attempting to describe. Adopting these criteria, risk assessments can to varying degree be judged as scientific in practice. Research on risk analysis, covering both foundational issues and practical implementation challenges, will contribute to improving the quality of risk assessment in this respect.

2 The risk concept

The term risk is defined in the literature in a number of ways. The presentation here is based on a distinction between overall qualitative definitions and their associated measurements/descriptions, and follows the one given by the new SRA glossary.

We consider a future activity [interpreted in a wide sense to also cover, for example, natural phenomena], for example the operation of a system, and define risk in relation to the consequences of this activity with respect to something that humans value. The consequences are often seen in relation to some reference values (planned values, objectives, etc.), and the focus is normally on negative, undesirable consequences. There is always at least one outcome that is considered as negative or undesirable.

Overall qualitative definitions of risk:

- a) the possibility of an unfortunate occurrence.
- b) the potential for realization of unwanted, negative consequences of an event
- c) exposure to a proposition (e.g. the occurrence of a loss) of which one is uncertain
- d) the consequences of the activity and associated uncertainties
- e) uncertainty about and severity of the consequences of an activity with respect to something that humans value
- f) the occurrences of some specified consequences of the activity and associated uncertainties
- g) the deviation from a reference value and associated uncertainties

These definitions express basically the same idea, adding the uncertainty dimension to events and consequences.

ISO defines risk as the effect of uncertainty on objectives. It is possible to interpret this definition in different ways; one as a special case of those considered above (e.g. d) or g)).

To describe or measure risk – to make judgments about how large or small the risk is, we use various metrics:

Risk metrics/descriptions (examples)

1. The combination of probability and magnitude/severity of consequences
2. The triplet (s_i, p_i, c_i) , where s_i is the i th scenario, p_i is the probability of that scenario, and c_i is the consequence of the i th scenario, $i = 1, 2, \dots, N$
3. The triplet (C', Q, K) , where C' is some specified consequences, Q a measure of uncertainty associated with C' (typically probability) and K the background knowledge that supports C' and Q (which includes a judgment of the strength of this knowledge)
4. Expected consequences (damage, loss), for example computed by:
 - i. Expected number of fatalities in a specific period of time or the expected number of fatalities per unit of exposure time
 - ii. The product of the probability of the hazard occurring and the probability that the relevant object is exposed given the hazard, and the expected damage given that the hazard occurs and the object is exposed to it (the last term is a vulnerability metric)
 - iii. Expected disutility.
5. A possibility distribution for the damage (for example a triangular possibility distribution)

The suitability of these metrics/descriptions depends on the situation. None of these examples can be viewed as risk itself, and the appropriateness of the metric/description can always be questioned. For example the expected consequences can be informative for large populations and individual risk, but not otherwise. For a specific decision situation, a selected set of metrics have to be determined meeting the need for decision support.

From this basis we are naturally led to interpretations of risk assessments, risk management, etc.

3 Risk management principles

Three major strategies are commonly used to manage risk: risk-informed, cautionary/precautionary and discursive strategies. In most cases the appropriate strategy would be a mixture of these three strategies.

The risk-informed strategy refers to the treatment of risk – avoidance, reduction, transfer and retention – using risk assessments in an absolute or relative way. The cautionary/precautionary strategy is also referred to as a strategy of robustness and resilience,

and highlights features like containment, the development of substitutes, safety factors, redundancy in designing safety devices, BAT (best available technology), as well as strengthening of the immune system, diversification of the means for approaching identical or similar ends, design of systems with flexible response options and the improvement of conditions for emergency management and system adaptation. A key aspect here is the ability to adequately read signals and precursors of serious events. All safety regulations are based on some level of such principles to meet the uncertainties, risks and the potential for surprises. The discursive strategy uses measures to build confidence and trustworthiness, through reduction of uncertainties and ambiguities, clarifications of facts, involvement of affected people, deliberation and accountability.

Risk assessment produces a risk description/characterization, which decision-makers and other stakeholders can use to support the decision-making and their views on relevant issues, such as choosing between alternatives, acceptance of activities and products, the implementation of risk-reducing measures, etc.

The risk assessment and its description/characterization must be seen as judgements made by the analyst group and the experts, that have been used to carry out the assessment. These judgments are conditional on a specific background knowledge which covers data, information and justified beliefs often formulated as assumptions. The risk-informed approach for managing risk, that originates from its assessment, has limitations in that it does not in general capture all aspects of concern for the decision making, nor all contributions to risk and uncertainties (as the description/characterization express conditional risk and are dependent on the analysts and experts conducting the assessments).

The generation of the risk-information could be supplemented with decision analysis tools such as cost-benefit analysis, cost-effectiveness analysis and multi-attribute analysis. All these methods have in common that they are systematic approaches for organising the pros and cons of a decision alternative, but they differ with respect to the extent to which one is willing to make the factors in the problem explicitly comparable. Independent of the tool, there is always a need for a managerial evaluation and review, which sees beyond the results of the analysis and add considerations linked to the knowledge and lack of knowledge that the assessments are based on, as well as issues not captured by the analysis. The degree of “completeness” of an analysis depends on the quality of the analysis and applied cut-offs.

The ALARP principle (ALARP: As Low As Reasonably Practicable) is a commonly adopted risk reduction principle, founded on both risk-informed and cautionary/precautionary thinking. The principle is based on the idea of gross disproportion and states that a risk-reducing measure shall be implemented unless it can be demonstrated that the costs are in gross disproportion to the benefits gained.

A challenge in risk management is to obtain sufficient focus on overall performance and be able to see both the upsides and downsides of the activity considered. Risk management is really about balancing the conflicts inherent in exploring opportunities, on the one hand, and avoiding losses, accidents and disasters, on the other. For some fixed frames it may be a goal to reduce risk (for example, risk related to accidents in a process plant), but it has to be acknowledged that the pursuit of some benefits could increase some risks over time. Different “scientific schools” have a special focus on these aspects – the need for seeing risk in relation to developments and overall performance, such as the quality discourse and Nassim Taleb’s antifragile conceptualization. Here the dynamic aspects of risk and

performance are highlighted, and it is stressed that to obtain top performance over time one has to acknowledge and even “love” some level of variation, uncertainty and risk.

There is a common belief among many – including analysts and managers - that to manage an activity, and avoid accidents and perform operations as planned, it is sufficient to develop procedures and ensure compliance with these. However, in practice, such a compliance perspective fails for non-trivial activities, since a “perfect” system cannot be developed; surprises always occur. The system understanding is too static, and improvements and excellence are not sufficiently stimulated. We have to acknowledge the performance, risk and knowledge “dynamics”. We need to see beyond compliance. For many type of systems, the signals and warnings are of a form that requires judgements and actions that need considerations beyond specified procedures.

4 Uncertainties in Risk analysis

When discussing uncertainties in a risk context, we need to clarify:

- i) What is uncertain?
- ii) Who is uncertain?
- iii) How should we represent the uncertainties?

In relation to i), we may be interested in uncertainties about an unknown quantity, uncertainties regarding what the consequences of an activity will be, and uncertainty related to a phenomenon, for example related to cause-effect relationships. When it comes to ii), we need to clarify if it is the decision-maker, the analyst or some experts used in the assessment, who are uncertain. In order to obtain a clear understanding of the risk and uncertainties and to communicate relevant results, it is essential to be precise on this issue. To express the uncertainties, an adequate representation is required, and probability is the natural choice as it meets some basic requirements for such a representation (Bedford and Cooke 2001, p. 20):

- *Axioms*: Specifying the formal properties of the uncertainty representation.
- *Interpretations*: Connecting the primitive terms in the axioms with observable phenomena.
- *Measurement procedures*: Providing, together with supplementary assumptions, practical methods for interpreting the axiom system.

A probability here means a subjective (judgemental) probability, expressing the assessor’s uncertainty (degree of belief) of the occurrence of event A. We denote this probability by P, or P(A|K) to show that this probability is conditional on some background knowledge K. A common interpretation is the uncertainty standard: the probability $P(A) = 0.1$ (say) means that the assessor compares his/her uncertainty (degree of belief) about the occurrence of the event A with the standard of drawing at random a specific ball from an urn that contains 10 balls.

Variation is also represented by probabilities, but of a different type, namely frequentist probabilities. A frequentist probability of an event A (denoted $P_f(A)$) is defined as the fraction of times the event A occurs if the situation considered were repeated (hypothetically) an infinite number of times. A frequentist probability $P_f(A)$ is thus a mind-constructed quantity;

a model concept. These probabilities constitute the basis for probability models, like the Poisson model to represent the variation of the number of events in specific periods of time.

Frequentist probabilities $P_f(A)$, thus, need to be justified. They can be constructed in cases of repeatability. When a frequentist probability can be justified, we may also refer to the propensity concept, that the probability exists per se; the probability is just a propensity of a repeatable experimental set-up, which produces outcomes with limiting relative frequency $P_f(A)$.

Normally a frequentist probability is unknown/uncertain. Uncertainties about its “true” value is then an issue. A subjective probability distribution can be used to express this uncertainty.

Frequentist probabilities model the variation in the phenomena, which is commonly referred to as aleatory or stochastic uncertainty. The subjective probabilities reflect epistemic uncertainties about the unknown quantity of interest.

The strength of the knowledge K supporting a subjective probability may be judged as more or less strong, depending on considerations of aspects like justification of assumptions made, amount of reliable and relevant data/information, agreement among experts and understanding of the phenomena involved. The set of probabilities and judgments of the strength of knowledge constitute together a broader way of expressing the uncertainties than just the probability assignments.

Interval probabilities are also used to express uncertainties. Instead of assigning a specific probability say 0.1 to an event, an interval say $[0.05, 0.3]$ is specified, meaning that the assigner is not willing to specify his/her uncertainty more precise than this. It does not mean that there is some uncertainty about where the “correct” or “true” probability is – there does not exist such a value - but there is imprecision in the sense that the assigner given the knowledge K is not willing to be more precise.

Such interval probabilities can be founded on alternative theories, such as possibility theory and evidence theory. By specifying a possibility function, such imprecision intervals will be generated. Also for such intervals, it is meaningful and relevant to consider the background knowledge and the strength of this knowledge. Normally the background knowledge in the case of intervals would be stronger than in the case of specific probability assignments, but they would be less informative in the sense of communicating the judgments of the experts making the assignments. Hence precise and imprecise probabilities can be viewed as supplementing measures rather than mutual excluding ones.

Models play an important role in risk analysis. To assess a quantity y we develop a model $g(x)$ as a function of some parameters x . The model error can be written $g(x)-y$ and is subject to uncertainty - we may refer to it as model uncertainty. Model uncertainty is also a key concept in case of phenomena (cause-effect) uncertainties. Think about the cause-effect relationships between smoking and lung cancer. Here we may develop a model of the number of deaths per (say) 100,000 persons (lung cancer mortality rate) in a population (for example women of a specific age group), linking it with the intensity (number of cigarettes per day) and duration of smoking (years) using a statistical model. Model uncertainty is uncertainty about the true difference between the model output and the actual mortality rate.

5 Confronting deep uncertainties, potential surprises and the unforeseen

A main issue in risk assessment and risk management is how to handle large/deep uncertainties in relation to the events occurring and the consequences of these events, such as in preparing for climate change and managing emerging diseases: what policies and decision-making schemes should be implemented in such cases? Traditional statistical methods and tools are not suitable, as relevant supporting models cannot easily be justified and relevant data are missing. However, other approaches and methods exist, such as methods for robust and adaptive risk analysis and management. These tools are based on two strategies; finding robust decisions that work acceptably well for many situations; and learning what to do by well-designed and analysed trial and error.

Adequately reading signals and warnings is a key tool in this respect. We must avoid missing or ignoring important early signals and precursors of serious events, but also exaggerating them. It is common practice to refer to false negatives (no indication of a risk situation when one is actually present) and false positives (erroneous signals indicating some risk situation is present when it is not), but how can we make judgements about these “errors”? It is easy to identify (and claim) that we missed a hazardous situation with hindsight a posteriori, when the accident, disaster or crisis has occurred, but how can we know in advance that we are missing, ignoring or exaggerating signals or precursors, given that we are typically exposed to a large number of threats/hazards? The reference for our evaluation of the signals and precursors cannot be the unknown consequences or outcomes of events yet to occur. The only possible way out seems to be to rely on the results of early risk and uncertainty assessments, where the warning system itself can be viewed as a form of risk assessment. However, risk assessment has its limitations as a tool for this purpose and, in cases when the knowledge base is not strong, we need to base the judgements on hypotheses and assumptions, and we may act too slowly (or too quickly).

This leads us to the use of adaptive risk analysis and robust analysis. Adaptive analysis is based on the acknowledgement that one best decision cannot be made but rather a set of alternatives should be dynamically tracked to gain information and knowledge about the effects of different courses of action. On an overarching level, the basic process is straightforward: one chooses an action based on broad considerations of risk and other aspects, monitors the effect, and adjusts the action based on the monitored results. A central idea is that since uncertainty is pervasive, one optimal management choice is not achievable. In this way, we may also avoid the occurrence of surprising type of events.

Abductive thinking is closely linked to adaptive analysis: you observe a surprising fact; in order to explain and understand this, you cast about in your mind for some theory or explanation. A new idea (or hypothesis) is brought up from the region where “all things swim”. In a process plant, abduction could mean to notice that the pressure is increasing, explore why and provide a hypothesis to explain it. Testing can, then, be carried out to prove or disprove the hypothesis. The approach is in line with fundamental ideas of the quality discourse, which highlights that knowledge is built on theory: rational prediction and analysis requires theory and builds knowledge through systematic revision and extension of theory based on a comparison of prediction with observation.

Bayesian decision analysis provides a strong theoretical framework for choosing optimal decisions in the case of information in the form of signals and warnings, but is difficult to use in practice, in particular in cases of large/deep uncertainties – the background knowledge is poor and the probabilities assigned may be difficult to justify. There are also well known challenges related to specifying utility functions to reflect the decision maker preferences.

For studying the robustness we can use many types of set-ups, for example (C,P,u,a), where C are the consequences of the actions, P the probabilities of C given the actions, u the utility function of C, and a the actions. However, it may be difficult to assign some of these values, for example P, when the uncertainties are large. A robust approach is, then, required. The key is to make decisions that are good for a set of values of, for example, C and P. However, the set-up used for robustness analysis may still exclude the possibility of many forms of surprises, as the framework reflects the current knowledge and beliefs. The protective measures could, for example, be based on a probability model reflecting variation due to a set of key risk sources but fail to include an important one. Robust analysis could be challenging to conduct in practice. There are many ways of looking at robustness and it may be difficult to find arguments for why some are better than others.

All this underlines the necessity for seeing the robustness analyses as decision support tools that need to be followed up with an evaluation and review, which also reflects on the choice of criteria or metrics used for the analyses, and the assumptions on which the analyses are based. The results of the analyses need to be seen in relation to these assumptions, as discussed above. Sensitivity analysis should always be used as well as some form of qualitative analysis of the arguments supporting different input values.

To meet deep uncertainties, potential surprises and the unforeseen, systems (organisations) must be made resilient; being able to sustain or restore their basic functionality following a stressor. A resilient system has the ability to

- respond to regular and irregular threats in a robust, yet flexible (adaptive) manner,
- to monitor what is going on, including its own performance,
- to anticipate risk events and opportunities,
- to learn from experience.

The focus on signals and precursors of serious events is a common feature of most approaches that intend to meet surprises and the unforeseen. If we look at basic insights from organisational theory and learning, we see that both this feature and resilience are main building blocks. A good example is the concept of collective mindfulness, linked to High Reliability Organizations (HROs), with its five principles: preoccupation with failure, reluctance to simplify, sensitivity to operations, commitment to resilience and deference to expertise. The thesis is that if organizations organize their efforts in line with these principles they will obtain high performance (high reliability) and effectively manage risks, the unforeseen and potential surprises.

Knowledge is the key to meet the potential surprises and the unforeseen. Risk assessment is about gaining, describing and communicating this knowledge, but the common approaches and methods used are not reflecting this dimension beyond assignments of probabilities and expected values. The current risk analysis practice mainly based on probabilistic thinking needs to be extended to meet the challenge of assessing and managing the risk related to potential surprises and the unforeseen. This is a huge research challenge for the risk field - it is further discussed in Section 7.

6 Reliability, validity and trustworthiness of risk analysis methods and results (including suggestions of how to make them more trustworthy)

In many cases, the evidence that the risk analysis is based on is very strong, and the message from the analysis is clear and trustworthy in the sense that people have confidence that the results describe the objective truth about something. If the results of the risk analysis present substantial amount of statistical data concerning the number of injuries and fatalities in relation to specific activities and populations, these data are not normally questioned as unreliable or not trustworthy. However, as soon as judgments are made about the future and cause-effect relationships, trustworthiness becomes an issue. Care must be shown in making statements about the true risks, in particular for cases characterised by large uncertainties. This was already discussed in Section 1: we need to make a shift in thinking, seeing risk analysis as a tool for revealing and describing the knowledge and the lack of knowledge we have about the phenomena studied. We should strive for faithful characterisations of the risk, uncertainties and knowledge, to support the relevant decision making.

To improve trustworthiness when adopting such a perspective, we have to make clear distinctions between data, information, know-how, justified beliefs and value judgments (that is, not free of controversial values).

Data are the symbolic representation of observable properties of the world, for example the failure data of a component or system. Information is processed data, for example mean failure times. Know-how is, for example, being able to accurately explain what a confidence interval means in a statistical setting and be able to conduct a statistical analysis according to the Bayesian procedure.

Risk assessments are full of “justified beliefs”, linked to probability judgments expressing the degree of beliefs of the assessor, for example about the strength of some material, and statements that are forming the knowledge basis of such probability judgments, for example, a belief that the system studied is of a standard type (and not of a different type with completely different characteristics). Here, the term ‘justified’ becomes critical. Philosophers and others have discussed the issue since ancient times.

The data and information provided by the various analyses and tests supply the premises for constructing a knowledge base, which is the collection of all “truths” (legitimate truth claims) and beliefs that the relevant scientists/analysts take as given in further research and analysis in the field. The evidence and the knowledge base are supposed to be free of non-epistemic values (“value-free” in the common but misleading jargon). Such values are presumed to be added only in the value stage, where conclusions are made on safety acceptance and too-high risk etc. Concluding that an activity is safe enough is a judgement based both on science/analysis and on values and norms. This is necessarily so, since the question “how safe is safe enough?” cannot be answered by purely scientific/analytical means.

The interpretation of the knowledge base is often quite complicated since it has to be performed against the background of general scientific knowledge. This is a step where the knowledge base is evaluated where one has to take the values of the decision-makers into account and a careful distinction has to be made between the scientific burden of proof - the amount of evidence required to treat an assertion as part of current scientific knowledge - and the practical burden of proof in a particular decision. However, the evaluation is so entwined with scientific issues that it nevertheless has to be performed or at least supported by scientific

experts. Many of the risk assessment reports emanating from various scientific and technical committees perform this function. These committees regularly operate in a “no man’s land” between science and policy, and it is no surprise that they often find themselves criticised on value-based grounds.

But it does not end there. In social practice, risk issues are often connected with – or rather are parts of – other issues. The decision-making has to be based not only on risk information but also on other concerns like costs. Even after the risk evaluation, decision-makers need to combine the risk information they have received with information from other sources and on other topics. There is a need for a decision-maker review and judgement which goes clearly beyond the scientific field and will cover value-based considerations of different types. It may include policy-related considerations on risk and safety, combining factual and value-based issues.

In its general form, risk is described by analysts by identifying a set of consequences of an activity and using a measure to express the uncertainties related to these consequences. This risk description is based on a background knowledge K , including assumptions on which the assessment is founded. Hence from the risk analyst perspective, risk is conditional given K . However, from the perspective of the decision maker an unconditional perspective is needed, seeing K as containing potential risk sources. To obtain trustworthy risk analyses, it is essential that this knowledge dimension of risk is understood.

7 The future of risk analysis: meeting the challenges. Emerging trends

Today risk analysis is well established in situations with considerable data and clear defined boundaries for its use. Statistical and probabilistic tools have been developed and provide useful decision support for many types of applications. However, risk decisions are, to an increasing extent, about situations characterized by large uncertainties and emergence. Such situations call for different types of approaches and methods, and it is a main challenge for the risk analysis field to develop suitable frameworks and tools for this purpose.

Known matters are the increasing complexity of coupled networks our society depends on, as well as those arising in preparing for climate change and in managing emerging diseases. Validated, trustworthy risk models providing the (aleatory) probabilities of the occurrences of future events and their consequences are not available, the relevance of past data for predicting future outcomes is in doubt, experts disagree — or reach an unwarranted consensus that replaces the acknowledgment of uncertainties and information gaps with group thinking — and policymakers are divided about what actions to take to reduce risks and increase benefits. Then, the question concerns what risk management policies to adopt in such cases. The use of robust and adaptive approaches can be seen as the answer, as discussed in Section 5, but there is a long way to go to equip the field with sound and practical methods for this purpose.

How shall we analyse and meet emergence risks - situations where the background knowledge is weak but contains indications/justified beliefs that a new type of event (new in the context of the activity considered) could occur in the future and potentially have severe consequences to something humans value? The weak background knowledge inter alia results in difficulty specifying consequences and possibly also in fully specifying the event itself; i.e., in difficulty specifying scenarios. We need risk assessments that are able to capture the

knowledge dimension and the time dynamics. A pure probabilistic approach – for example a Bayesian analysis, would not be feasible as the background knowledge - the basis for the probability models and assignments would be poor. For risk management, there is a need of balancing different risk management strategies in an adaptive manner, including (pre-)cautionary strategies and attention to signals and warnings.

Probabilistic risk assessment has proved to be a useful tool in a wide range of applications, but to meet the challenges of large uncertainties and emergent risks the probability-based approaches for assessing risk is not adequate. The argumentation follows different lines of thinking, but the main point being made is that the knowledge and information (or lack of knowledge and information) available for the analysis cannot be reflected properly by probabilities.

Approaches other than purely probabilistic have been suggested, for example using interval probabilities, possibilities, or qualitative methods, but there is no consensus on how to best meet this challenge. There is clearly a need for research in the development of sound approaches that can complement the traditional probabilistic risk analyses in appropriately reflecting uncertainties depending on the information available in support of the assessment and for the decision-making.

There is a need for substantial research and development to obtain adequate modeling and analysis methods – beyond the “traditional ones - to “handle” different types of systems. Examples include critical infrastructures (e.g. electrical grids, transportation networks etc.), which are complex systems and often interdependent, i.e. “systems of systems”. A comprehensive risk and vulnerability analysis of such systems requires the assessment of a broad spectrum of hazards and threats of different nature, with considerations of a large number of spatially distributed, interacting elements with nonlinear behavior and feedback loops. Another example is security-type applications, where qualitative assessments are often performed on the basis of judgments of actors’ intentions and capacities, without reference to a probability scale. There seem to be a huge potential for significant improvements in the way security is assessed by developing frameworks that integrate the standard security approaches and ways of assessing and threatening uncertainty.

Societal risk decision-making is more and more challenging – it is characterized by many and diverse stakeholders, and thus requires the integration of the diverse scientific, economic, social and cultural aspects which are represented, with the considerations of risk and uncertainties. The processes used for the integration include a variety of multifaceted, multi-actor risk exercises, but also need to include the consideration of other contextual factors such as institutional arrangements (e.g., the regulatory and legal frameworks that determine the relationships, roles and responsibilities of the actors), co-ordination mechanisms (such as markets, incentives or self-imposed norms) and political cultures, including different perceptions of risk. Some of the challenges and research issues that need to be focused on, here, relate to, inter alia:

- how the outcomes of the risk and uncertainties assessment should be best described, visualized and communicated, for their informative use in the above described process of societal decision-making involving multiple and diverse stakeholders,
- how issues of risk acceptability need to be seen in relation to the measurement tools used to make judgments about risk acceptability, accounting for the value generating processes at the societal level,

- how the managerial review and judgment should be defined in this context.

Key issues that we need to address are:

- In intergenerational decision-making situations, what are the available frameworks and perspectives to be taken? What are other options? When are different frameworks more appropriate than others? How do we capture the key knowledge issues and uncertainties of the present and future? What duty of care do we owe to future generations?
- How can we describe and represent the results of risk assessment in a way useful to decision makers, which clearly presents the assumptions made and their justification with respect to the knowledge which the assessment is based upon?
- How can we display risk information without misrepresenting what we know and do not know?
- How can we accurately represent and account for uncertainties in a way to properly justify confidence in the risk results?
- How can we state how good are expert judgments and how can we improve them?
- In the analysis of near-misses, how should we structure the multi-dimensional space of causal proximity among different scenarios in order to measure “how near is a miss to an actual accident”?

The above list covers issues ranging from important features of risk assessment to overall aspects concerning risk management and governance. Clearly, the list could be made much longer; it has to be seen as examples of issues that need further attention.