Norsk olje&gass

BLACK SWANS

An enhanced perspective on understanding, assessing and managing risk

This report is a translation of the Norwegian original prepared by Norwegian Oil and Gas Association in the project on “Black Swans - an enhanced perspective on risk”.
Project leader: Bodil Sophia Krohn

Stavanger, 25. August 2017

PREFACE

We use the concept of "black swans" to sharpen awareness of and caution over the uncertainty and risk which will always be associated with activities in the oil and gas industry. Risk needs to be understood and managed. Black swans will help us to detect uncertainty and risk before the threat of an incident and an accident becomes reality.

A Black swan is a metaphor used to obtain focus on the fact that serious events may occur - events we had not thought of or had knowledge about, had ignored or had considered not realistically happening. Black swans always take us by surprise in relation to our beliefs and knowledge.

This report will help us to achieve a better understanding of uncertainty and risk. It does not contain all the answers, but identifies directions for the work which can be done and provides a number of examples of how improvements can be achieved. We owe a big debt of gratitude to Professor Terje Aven at the University of Stavanger, who has been a key contributor to this report. Everyone concerned with safety in the oil and gas industry should benefit from familiarising themselves with its contents.

As an industry, we have achieved much. But we will never cease to need constant improvement. We will and must always stretch ourselves to achieve the best. We must detect the black swans in time.

CONTENTS

INTRODUCTION

This report presents the Norwegian Oil and Gas project on “-Black swans – an enhanced perspective on risk”. The purpose of the report is to explain what an enhanced risk perspective comprises, what can be achieved, and when it can be applied.

In an enhanced perspective, knowledge and uncertainty become important aspects of risk. That involves new and improved ways of viewing risk in relation to knowledge-building, experience transfer and learning. Drawing on disciplines which are particularly concerned with variation, knowledge and learning, provides new insights on ways of meeting the unthinkable, unforeseen and potential surprises. In particular, it demonstrates how ideas from organisation theory and lessons – rooted in the concept of “collective mindfulness” used in studies of high-reliability organisations (HROs) – can be related to broader thinking about risk. An enhanced perspective provides the basis for a more dynamic risk understanding which fits with the principles for collective mindfulness –for example by directing attention at signals, for example, and by being sensitive to what actually happens in a process. A dynamic understanding of risk is essential for dealing with risk in a complex system.

A potential exists for adjusting and simplifying many of the principles and methods applied, so that risk analyses and management can provide more appropriate decision support. Improved risk management can be achieved while also making safety work more cost-efficient.

The report is divided into four chapters. Chapter 1 provides an introduction to the enhanced risk perspective, what it means, what it provides and when it can be applied. Chapter 2 then reviews key concepts, such as risk, probability, knowledge and black swans. Chapter 3 discusses various topics related to analysing and managing risk in light of the enhanced risk perspective. Chapter 4 discusses some examples of black swans and actual incidents. The appendix provides a checklist of important aspects to look out for in order to take better account of the knowledge dimension and the unforeseen in connection with risk analyses.

1 AN ENHANCED PERSPECTIVE ON RISK

WHAT IS AN ENHANCED PERSPECTIVE ON RISK ABOUT?

WHY APPLY AN ENHANCED PERSPECTIVE ON RISK?

WHEN CAN AN ENHANCED PERSPECTIVE ON RISK BE USED?

1.1 What is an enhanced perspective on risk about?

An enhanced perspective on risk builds on established principles and methods. The bowtie principle, for example, is an excellent tool for identifying causes and consequences of undesirable incidents and for assessing barriers. Many other principles and methods are also available for understanding and analysing risk. But something more is needed.

Risk cannot be described only with the aid of historical figures, average values and probabilities. The risk matrix alone, with its basis in probabilities and consequences, can provide a misleading picture of risk. We need to look beyond probabilities.

An enhanced risk perspective provides a broader basis for understanding risk. Our knowledge and lack of it (uncertainties) about the phenomena we study are as important as the probabilities we arrive at. Moreover, we must be aware that surprises can occur with the assessments we make. Ways of thinking (mindset) and argumentation could be restricted. Important risk aspects could lie concealed here. An enhanced risk perspective covers all these considerations.

With an enhanced risk understanding of this kind, greater emphasis than before is placed on the significance of knowledge-building, experience transfer and learning. Knowledge forms part of the risk description.

However, the enhanced perspective on risk involves rather more than this. In particular, it covers concepts and instruments related to understanding, assessing and managing the two new risk components – the knowledge dimension and the unforeseen/surprises.
See figure 1.

Figure 2 points to two new areas from the quality discipline and organisational theory, which refers here to the collective mindfulness concept used in studies of HROs. The quality discipline emphasises knowledge, system understanding and insight by studying variation, at the same as surprises, learning and continuous improvement are constantly considered.

The collective mindfulness approach involves five main principles:

- Preoccupation with failure.
- Reluctance to simplify.
- Sensitivity to operations.
- Commitment to resilience.
- Deference to expertise.

When assessing and dealing with risk related to an activity, these principles can be used as an aid. If, for example, risk is assessed on the basis of a greatly simplified risk description (such as a risk matrix), the principle of reluctance to simplify would be breached. In an operation, the ability to pick up signals and to make adjustments or corrections will often be crucial for achieving a desired outcome (sensitivity to operations, preoccupation with failure). Variation is measured with the aid of various indicators. The challenge is to distinguish the special

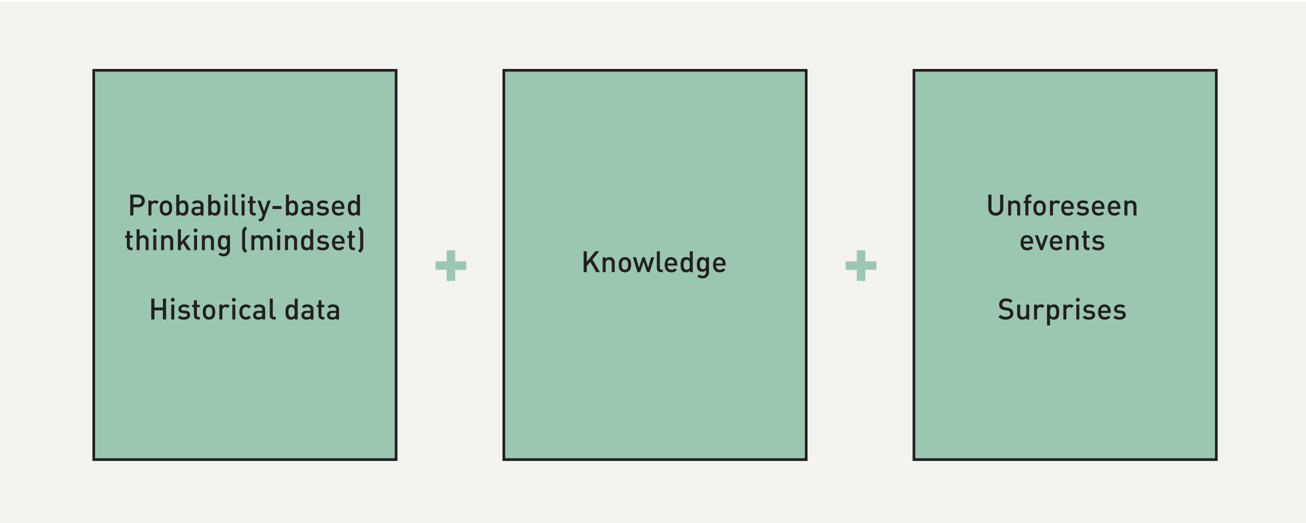

FIGURE 1: WHAT RISK COVERS

variation, which requires immediate follow-up, from the normal variation. The quality discipline provides concepts and insights that are important in this regard.

Unforeseen incidents and surprises largely signify our failure to see what is coming, even when we ought to. Our concern in an enhanced perspective is not only with avoiding incidents but also with developing good and ever-better solutions.

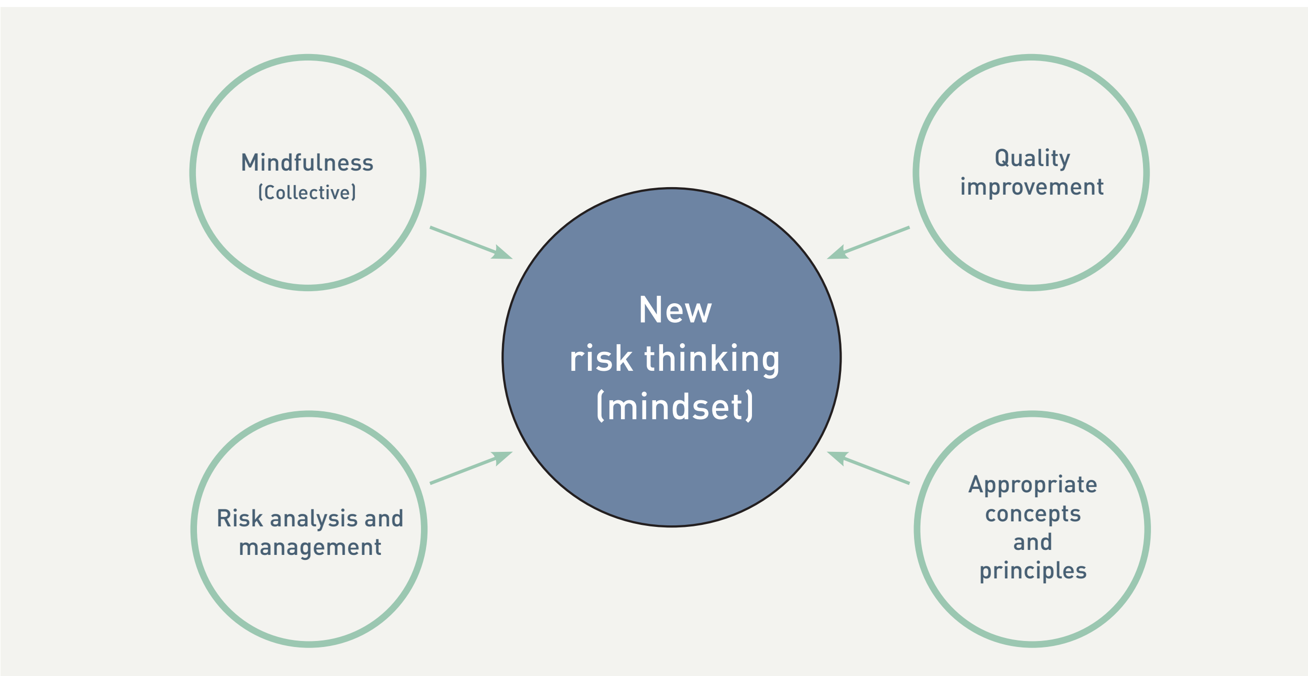

FIGURE 2: IMPORTANT PILLARS OF THE NEW RISK THINKING (MINDSET)

1.2 Why an enhanced perspective on risk?

The risk thinking (mindset) has traditionally been based to a great extent on historical data and probabilities – through the use of risk matrices, for example, which express the probability of given events (incidents) and typical consequences of these events if they occur. The most important arguments for an enhanced approach to risk are presented in figure 3, and elaborated in the following.

- **Excessive emphasis on probabilities and historical data.** Risk is about the future, and that can be characterised to a greater or lesser extent by what has happened in the past. History can not capture all the types of events which could occur. Probabilities can be used to describe variation, but it is often difficult to justify a specific probability model to describe this variation. Probabilities can also be used to express beliefs how likely it is that an event will occur, but this too has its clear limitations.

  This is described in the following points:

  The concentration on probabilities and historical data has a tendency to coincide with an understanding of risk being an inherent property of the system that cannot be easily influenced. This is an unfortunate understanding, which breaks with intuition and with the reality people experience when they are involved. See also the point below on signals.

- **Probabilities may be the same, but the knowledge they build on may be strongly or weakly founded.** For examples in two different cases, a probability of 0.50 could be specified – getting heads when tossing a coin or a possible terrorist attack. In the first case, the knowledge on which the percentage builds is very strong. In the other, it could be extremely weak. Looking solely at probability figures could mean that important information gets overlooked. See also chapter 2.4.

- **Important aspects of risk and uncertainty are concealed. The assumptions may conceal important uncertainty factors.**
  A risk statement on the basis of probabilities builds on assumptions and background knowledge which could hide significant contributors to risk. Risk analyses are always founded on assumptions. It might be assumed, for example, that a specific procedure is followed, but this can never be certain in practice. A risk of assumption deviation is present, which is not sufficiently addressed today.

- **Dominant explanations and mindsets may prove incorrect.** The methods and models used are only aids. A humble attitude is needed over what these methods and models can and cannot provide. Unforeseen events and surprises happen in relation to common and established ideas and mindsets. So a conceptual framework, a mindset and methods which reflect this are required.

- Excessive emphasis on probabilities and historical data
- Important aspects of risk and uncertainty are not highlighted
- The assumptions may conceal substantial uncertainty factors
- Dominant explanations and beliefs (mindsets) could turn out to be wrong
- Probabilities may be the same, but the knowledge they build on may be strongly or weakly founded
- Many signals may exist, but these are not reflected in the probabilities and the failure data
- Considerable variations between different groups in the way risk is understood
- Variation is seen as something negative

FIGURE 3: CHALLENGES WITH THE EXISTING RISK THINKING (MINDSETS)

- **Many signals may exist, but these are not reflected in the probabilities and the failure data.** If the risk figures are established solely on the basis of observed failures and accidents, the risk picture will be oversimplified. We are unable to reflect the information and knowledge available in a given context.

- **Considerable variations between different groups in the way risk is understood.** Today's reality is characterised by many divergent views and perspectives about understanding of what is risk and how it should be described. Many of these are unfortunate – regarding risk solely as an expected value (where probability and loss are multiplied), for example.

- **Variation is regarded as something negative.** We want to meet requirements, but the attention paid to this may become so strong that it can obstruct good improvement processes. Deviations can lead to substantial growth and progress in the right setting and culture. We must seek to develop good solutions in order to improve. This concerns more than a sole concentration on measurements to avoid nonconformities.

WHAT BENEFITS DOES AN ENHANCED PERSPECTIVE ON RISK PROVIDE?

- **Improved risk analyses**
  The ability to perceive the overall picture, and to reflect knowledge and the lack of it. Be alert and pick up nuances and signals when something happens. Take account of the unforeseen and think resilience (a form of robustness – see section 2.5).

- **Improved risk analyses**
  Capture and describe this risk.

- **Improved communication**
  Communicate the understanding of risk

- **Improved risk management**
  Use the enhanced risk understanding and analyses as a basis for the decisions to be taken in order to choose the right solutions and measures.

  Prevent a major accident from taking place.

1.3 When can an enhanced perspective on risk be used?

An enhanced perspective on risk can be applied in all circumstances where risk is faced.

This is a matter of possessing a good understanding of risk, and drawing on it to adopt the right actions and measures.

That could apply to a team planning a set of drilling and well operations or to personnel preparing to undertake maintenance work – on hydrocarbon-bearing pipes and equipment, for example. Everyone faces the fact that undesirable events can occur, and that these could have serious consequences. Before the activity, we do not know what events that will occur and what they will lead to. Risk is present.

An enhanced risk perspective is relevant in all phases of a project or activity, at the planning stage as well as during operation and at cessation. Risk thinking also affects us when an undesirable event has happened, such as a gas leak. That influence is exerted through the understanding of risk and the action taken to meet such events, and the way the situation is dealt with.

2 CENTRAL CONCEPTS

2.1 What is risk?

Risk relates to accidents, terrorism, illnesses, financial loss and war. But what does it actually mean? Formulating this precisely is not so easy. Many people have problems explaining risk precisely. It is an abstract and difficult concept.

Many people think in terms of what has happened –events, losses and crises. But these are not risk. They are observations and history. No physical law says that what has happened will repeat itself. We can learn from it, and we can use history to say something about what we think the future might hold. But the future will always be uncertain – whether incidents will happen or not, and what their consequences (outcomes) will be if they actually do occur. How many will be injured or killed, how big the financial losses will be and so forth? This is risk in our context – that incidents with different consequences could happen. Risk accordingly has two main components: (i) the event and its associated consequences and (ii) uncertainty about these – will the incidents occur and what will their consequences be? Taken together, these two components constitute risk.

The size of the risk then becomes the interesting question. Risk analyses are carried out and the risk is described. Probabilities then enter the picture, and it is possible to express whether the risk is large or small. Unfortunately, however, looking at the probabilities is not enough for assessing whether the risk is high or low. These figures are only a tool for expressing risk. Account must also be taken of the available knowledge and its strength. See the discussion in the following sessions.

Reference is made in the report to the new definition of risk provided in the guidelines to section 11 of the framework HSE regulations: *"Risk means the consequences of the activities, with associated uncertainty".* This definition of risk corresponds to a great extent to the definition of risk in ISO 31000 (2009, 2.1): "risk is the effect of uncertainty on objectives".

2.2 What is a probability?

A probability provides a way of expressing uncertainty – or, put another way, the degree of belief of a specific event occurring. If the probability of an event is 10 %, it implies that one judges its probability to be the same as drawing a specific ball at random from an urn containing 10 balls.

AN EXAMPLE

Let A be the probability of a large leak occurring on offshore installation B next year ("large" is precisely defined, but this is not important here). At present, we do not know whether A will occur or not. Uncertainty is related to the outcome. We can have a view about the likelihood of such an event occurs. We can look at all the available statistics on such leaks, which perhaps indicate a historical rate of 2/100 per annum ¬– subject to certain conditions. The relevant installation has not experienced a leak of this kind before, but quite a lot of smaller ones. Identifying which installations we should use for comparative purposes, and which types of incidents to take into account, is obviously a challenge. Other assumptions generate different figures.

In order to determine the probability of A occurring, we must take a view on these data and assess their relevance to the installation concerned and to next year's operation. We may conclude with a probability of 5/100. This expresses how likely the analyst believes event A to be, given their background knowledge (justified beliefs) which is based on data, information, modelling, argumentation, testing etc., and is often formulated as assumptions in the assessments. We call this background knowledge K, and can then express the probability of A as P(A|K), where | indicates that the probability is given by (conditional on) K. This is often simplified to P(A), with the background knowledge K taken as read. We refer to this as a knowledge-based probability, or simply a probability.

If we say the probability is 5/100, we believe that it is equally likely that the event A occurs as drawing a red ball from an urn containing five red and 95 differently coloured balls. The uncertainty is the same. We see that we can also understand a probability as an expression of uncertainty about the outcome. However, a simpler approach is to think of it as expression for the likelihood of the event occurring (degree of belief). As a rule, therefore, we will refer to this interpretation. To put this another way, if we can imagine 100 installation-years which are relatively similar to the next one, we will predict five such leaks over that period. No correct/true/objective probability exists here. No correct probability exists even if you roll a dice. That may sound strange, but we need to distinguish between proportions, observed or imagined, and probability as used here.

A DICE EXPERIMENT

If we imagine a dice and roll it a large number of times – let us say 6 000 – we end up (if the dice is fair) with about 1 000 ones, 1 000 twos and so forth. In a population of 6 000 rolls, the distribution will be pretty similar to 1/6 for each of the different numbers. We can also conceive making an infinite number of rolls. The theory then says that the distribution will be exactly one-sixth. But these are proportions, observed in and the results of an imagined experiment. They are not probabilities which express degrees of belief and the analyst's uncertainty about a defined event which we do not know will happen or not. If we roll a dice, it can show a four or another number. Before rolling, we can express our belief that the dice will come up four. We will usually specify this probability as 1/6, on the basis that all outcomes must be equally probable because of symmetry (we are talking about a classic probability expressed as the number of favourable outcomes divided by the number of possible outcomes). This will clearly give the best prediction of the number of fours if we roll many times.

With a standard dice, we expect to obtain a four about 1 every 6th rolls in the long run. But there is nothing automatic about setting the probability at 1/6. We must make a choice. It is us to determine the likelihood to achieve a four, given our background knowledge. If the dice is known to be fair, 1/6 is the natural choice. We may be convinced that the dice is not standard and will yield far more fours than normal. The assignment could then be P(four) = 0.2. Nobody can say this is wrong, even if a subsequent check of the proportion of fours for this dice shows it to be standard. The background knowledge was different when the probability was determined. The probability must always be viewed in relation to its background knowledge.

Classical statistics builds on a completely different understanding of what probability is. It is defined there as the limit of a relative frequency (we call this a frequentist probability) – in other words, the proportion outlined above when the number of trials becomes infinite. Faced with a new type of dice, a frequentist probability p of the event "four" is defined by imagining this trial conducted an infinite number of times. That approach establishes "true"/"objective" probabilities, with efforts made to estimate these through trials and analyses. These frequentist probabilities are perceived to be a property of the phenomenon under study.
Care needs to be taken in applying this understanding of probabilities to the real world.
It is easy to be misled into thinking that certain events "must" happen – a kind of fatalism. Since such thinking ascribes a frequentist probability to all types of events, the latter will occur in accordance with that probability. In reality, however, stable processes generating such frequentist probabilities do not exist. They occur only in an imaginary world built up with the aid of models and a number of assumptions. The actual outcome for the relevant installation depends on what actually happens out there in terms of humans, technology and organisation.

No physical law specifies that a large leak must occur every nth year. In fact, it may never happen. Knowledge-based probabilities express the analysts' degree of beliefs that certain events will happen. That represents something completely different. It does not claim to be a description of the way the world behaves, but of how probable something is from the analysts' perspective.

If we say that the probability of an event is consider to be negligible, we have strong belief that it will not happen. Our knowledge indicates that this will not occur. We are pretty sure of that. Should the event nevertheless happen, that would be a surprise relative to our knowledge/beliefs. It is important to remember that the probability of an event viewed in isolation could be low, but would increase substantially for a larger population. A probability of 1/10 000 could have been assigned for a certain event on a specific installations, for example. If we look at 100 installations, however, the probability would be 100 times greater (in other words, 1/100) and the event would not be so surprising if it occurs.

2.3 What are knowledge and uncertainty in a risk context?

The classical definition refers to three criteria for a claim to be considered knowledge: it must be true, the individual must believe it to be true, and it is justified. In our setting, knowledge is justified beliefs – we must avoid all references to truth for the definition to function in our setting.

When we say risk assessments are conditional on a given background knowledge, that must be interpreted on the basis of this understanding of knowledge. However, it could be appropriate – in line with the normal classification of knowledge – to distinguish between data (D), information (I), knowledge (K) in the sense of justified beliefs, and wisdom (W) – in other words, the four elements in the DIKW hierarchy.

Data D could be observed operational lifetimes for a type of unit, information I could be estimated failure rates derived from D, knowledge K could be the probability model utilised and wisdom W could be proper use of results from the risk analyses, where adequate weight is given, for example, to the limitations and assumptions of the analysis.

We distinguish between uncertainty about an unknown quantity, what the consequences of an activity will be, and in relation to a phenomenon such as the climate. The challenge is to conceptualise and measure uncertainty. The basic thesis is that uncertainty about a quantity or the consequences of an activity involves not knowing either the value of this quantity or the consequences. Measuring this uncertainty leads to concepts such as probability, with the addition of the background knowledge on which this measurement is based. Where uncertainty about a phenomenon is concerned, concepts such as scientific uncertainty and lack of (good) predictive models are used. A normal understanding of scientific uncertainty is that no scientific consensus exists on a prediction model f which permits the quantity we want to study, X, to be expressed as $X = f(Z1, Z2 \ldots)$ with a high degree of confidence, where Z1, Z2 ... are underlying factors which influence X.

Uncertainty can be defined and classified in a number of other ways. One example is the division into unknown unknowns, unknown knowns and so forth. See sections 2.6 and 2.7. An important point here is to clarify uncertainty and lack of knowledge for whom and when.

2.4 Risk descriptions and probabilities depend on the strength of the knowledge these are build on

A probability expresses the degree of belief that an event will occur given a specific background knowledge, as discussed in section 2.2. Let us assume that a probability of 0.50 is assigned in a specific case. This value could build on strong or weak background knowledge, in the sense that considerably more relevant data and/or other information and knowledge which supports a value of 0.50 is available in one case than in another case with few or no data or other information/knowledge to support such a value. Let us consider an extreme case. You have a normal coin in front of you and toss it once. You specify a probability of 0.50 for heads – your background knowledge is obviously strong. The specification is based on the reasoning that both sides are equally probable because of symmetry, and experience from such coins supports the view that you will get roughly 50% heads or tails.

But let us now imagine that you are going to determine a probability of heads for a new coin you know nothing about. It could be normal or abnormal (you are not allowed to see it). What will your probability be? Most people would probably say 50 % again. But your background knowledge is now weak. You have little insight into what kind of coin this is. We see that the probability remains the same, but the background knowledge is strong in one case and weak in the other. When assessing the "power" of an assigned probability, it is obviously important to learn more about this background knowledge. The figure alone does not reveal that much. This will also be the case when we use probabilities to describe risk. The figures build on some background knowledge, and we must know its strength in order to use them correctly in the risk management. What about all the assumptions made? How reliable are they? How good are the underlying data and models? What about the expert assessments incorporated? Do the various experts agree? How well do we understand the phenomena involved? Has the knowledge which underpins the probability assessments been subject to a thorough review?

The question of the strength of background knowledge will also be relevant where qualitative risk analyses have been carried out without probability figures. All risk assessments are based on knowledge (through data, models, expert judgements..) and the strength of this knowledge must always be considered for the risk description to be meaningfully interpreted and adequately support the decision-making.

2.5 What do vulnerability, robustness and resilience mean?

If we say that a person who contracts an infectious disease is vulnerable to it, this could mean that he/she has a high "risk of death" as a result of the infection. Vulnerability can accordingly be perceived as conditional risk, given an event, stress or a form of exposure. In other words, everything we have said above about risk can be applied to vulnerability, providing we take care to include this condition.

Let us look at some other examples. The degree of vulnerability can be expressed in part as:

- the probability that a desired function – such as power supply – cannot be maintained given an initiating event.
- the probability of a person in a car being seriously injured given a defined collision scenario.

In its 2000 report (NOU 2000:24), the Norwegian government commission on the vulnerability of society defined vulnerability as follows:

Vulnerability is an expression of the problems a system will have to face when subjected to an adverse event, as well as the problems for the system in resuming its activities after the event has occurred. Vulnerability is related to the possible loss of value. A system in this context is, for example , a state, national power supply, a company or an individual computer system.

All these examples fit with our general understanding of vulnerability. Given the event (or risk source) A, some consequences (relative to the values we are concerned with) could be the result. We do not know which. This is the uncertainty.

These consequences are studied in risk analyses. We use probability to express the uncertainties and how likely this or that consequence will be. The knowledge on which these assessments build are an integrated part of this conditional risk description assuming the event occurs, exactly as is the case with the unconditional risk description.

When we talk about a system being vulnerable, we mean that vulnerability is assessed to be high. The point is that we regard the risk to be high, assuming that the system experiences an initiating event. If we know that the system lacks effective barriers or emergency response systems should a fault occur, we would say that the vulnerability is high.

ROBUSTNESS

The term "robustness" is used in various ways. In this document, we mean the opposite of vulnerability in the sense that low vulnerability means a high level of robustness and vice versa.

RESILIENCE

Resilience is closely related to vulnerability and robustness. The term refers to the ability of a unit (organisation) to recognise, adapt to and absorb variation, change, disruption and surprise. In other words, the term goes beyond robustness in that we are not starting from a defined event (risk source). We are concerned with the way the system will function if it is also exposed to a source or incident we have not previously thought about. Measuring and describing the degree of resilience are accordingly difficult. How can we measure the way our body might cope with illnesses we know nothing about today? The concept is nevertheless important because resilient systems are clearly desirable – they can cope with both known and unknown stresses (changes).

A lecturer is resilient in a certain sense if he/she succeeds in tackling unforeseen questions and comments from the audience. We can measure this ability by looking at how the person concerned has dealt with unforeseen questions before, and we can conduct trial lectures, but the limitations of such an approach quickly become apparent. The lecturer may have several weak points – we can call them vulnerabilities in relation to very specific types of questions – but their content could be unknown to everyone. As already noted, however, this does not mean the concept is not useful. We face the same challenge with risk and the ability to detect surprises.

2.6 What is a black swan in a risk context?

A black swan in this context is understood as an event (or a combination of events and conditions) which is unforeseen and/or represents a surprise in relation to our knowledge and beliefs. When speaking about such events, the implication is always that their consequences are serious.

A LIQUID EXAMPLE

Think of a container for liquid, which is normally filled with water and which people drink from every day. On one occasion, John drinks a liquid from the container which proves to be poisonous. This incident is a black swan for John, a surprise in relation to his knowledge and beliefs (and has serious consequences). Afterwards, an explanation is found for why the liquid was toxic. That is also characteristic of a black swan – it can be explained with hindsight.

A specific example of a black swan of this kind, based on Sande (2013), relates to open-air explosions. Until around 1960, it was accepted knowledge that natural gas cannot explode in the open air. This knowledge was supported by experiments, but the latter used inadequate gas volumes. Later trials with larger quantities of various gas types have shown that all hydrocarbon gases can explode in the open air. However, different gases vary in terms of their "critical diameter" – the smallest diameter which must be present for a continuous detonation to occur.

THREE TYPES OF BLACK SWANS

Three types of black swans are defined in our context. See figure 4. The first is the extreme – the completely unthinkable– an event unknown to the whole scientific community, to everyone. A new type of virus provides an example. This type of event – often termed unknown unknowns – cannot be excluded, but is not so "likely" where our knowledge is strong and our experience extensive, as in the petroleum industry.

The second type of black swans is events that we have not included in our risk assessment, either because we do not know them or because we have not made a conscious consideration. The poisonous liquid in the example can be viewed as a black swan on this basis. John's risk analysis was pretty superficial, if he carried one out at all. Regardless, the poison took him by surprise because he had not imagined that such a scenario was a possibility. In other words, it was unforeseen.

If John had conducted a more detailed risk analysis, such a scenario would have been identified because people know it can happen and everyone has heard about it. Nevertheless, the incident could have been a surprise because the risk assessment had indicated that the probability was so small that it could be disregarded. That represents the third type of black swan – incidents which occur even though we have assessed their probability as negligible.
Let us modify the liquid example a little. Assume

that a risk analysis has identified various types of poisonous liquids which could fill the container in special circumstances, but excludes a hazardous form because of a set of physical arguments. But this scenario occurs nevertheless. It happens despite being regarded as impossible by the analyst. The real conditions were not the same as those applied in the risk analysis, and the incident took even the analysts by surprise. However, it was easy to explain with hindsight. This incident was consequently in the third category

We could also have had a black-swan event of the second category. Limited knowledge of the subject causes the risk analysts to overlook a certain type of liquid (known to other specialists) in their analysis. Its appearance accordingly becomes a black swan for the analysts. We can call this an "unknown known", known to some, unknown to others.

An important form of black swan is the combination of events and conditions which leads to disaster. Such combinations are precisely what characterises serious accidents. Risk analyses fail to detect them and, even if they are spotted, and would normally ignore them because of their negligible probability.

Look at the 2010 Deepwater Horizon disaster in the Gulf of Mexico. A blowout is not a surprising event in itself – it is an obvious event for risk analyses – but this event involved an unforeseen combination of events and conditions (PSA 2013):

- Erroneous assessments of pressure test results.
- Failure to identify formation fluid penetrating the well despite log data showing that this was happening.
- The diverter system was unable to divert gas.
- The cutting valve (blind shear ram – BSR) in the blowout precentor (BOP) failed to seal the well.

This was probably not a relevant scenario ahead of the accident . But it nevertheless happened.

PERFECT STORMS

Many readers will have seen The Perfect Storm, the American film starring George Clooney. It deals with a terrible storm in the Atlantic which killed 12 people, and which arose from the combination of a storm starting over the USA, a cold front coming from the north and the remnants of a tropical storm originating to the south. All three types of weather are known from before and occur regularly, but the combination is very rare. Captain Billy Tyne (Clooney) and his crew decided to take the risk and ride out the storm – but had not foreseen its strength. Their boat is destroyed and nobody survives.

This extreme event is now used as a metaphor for a rare incident which can occur and where we understand the phenomena involved. Experts can calculate their probability and the associated risk very precisely, and can largely predict the course of events. They can say that one in 10 of such events will involve waves so high, 1 in 100 will have waves of that height and so forth. Such events are taken into account when building offshore installations, with strength standards to ensure that they can withstand extreme waves. But there is always a limit. We must accept that one wave could be so high that the platform fails to survive it, although such an incident has a very low probability.

Similarities to that position can be found in other areas, such as health care and road traffic. We know fairly precisely what proportion of the population will suffer a given illness over the coming year and how many will die on the roads. Measures can be taken to reduce the risk, and we can measure possible changes over time. The annual number of people killed in road accidents since the 1970s has declined steadily although traffic has risen. Risk management unquestionably works.

“The Perfect Storm” (2000)

In other words, the “perfect storm” metaphor relates to events where traditional science prevails, where we have precise probabilities and relevant statistics, and where we can make accurate predictions about the future.

The rare incident in the perfect storm corresponds ostensibly to type c) of black swans, but there is an important distinction. The variation in phenomena in perfect storms is known – uncertainty is low, the knowledge base is strong and accurate predictions can be made. When the knowledge base is so strong, black swans can be ignored for all practical purposes. The probabilities are frequentist and characterise variations in the phenomenon being studied, and are known to an extent regarded as certain.

Where black swans of type c) are concerned, the position is generally that this kind of accurate prediction cannot be made. Variation in the phenomenon cannot be described with such precision. We must rely on assessments and knowledge-based probabilities which express degrees of beliefs for the events to occur. Such a perspective is applied when the probability of an event is said to be negligible and is accordingly not believed to happen. We can obviously experience surprises in relation to such assessments.

2.7 What is meant by saying an incident is unthinkable, unforeseen or surprising?

We use the terms unthinkable, unforeseen and surprising interchangeably and with the same content as the black swan discussed above when we restrict ourselves to events where serious consequences are implicit. In other words, these terms cover:

a. unknown unknowns
   – events unknown to everyone
b. unknown knowns
   – events we have not included in our risk assessment, either because we are not aware of them or because we have not conducted a conscious assessment, but which are known to others
c. known events, but which are assessed to have a negligible probability, and which we accordingly do not think will happen.

However, some refinements are appropriate. See figure 4.

An event in categories b) and c) will clearly take us by surprise, but this is less obvious for unknown unknowns (category a). After all, we may have an activity where considerable uncertainty prevails about the types of events and consequences which it will generate, and we take an open and cautious attitude on what could happen. We have no specific expectations. Whether an unknown unknown which

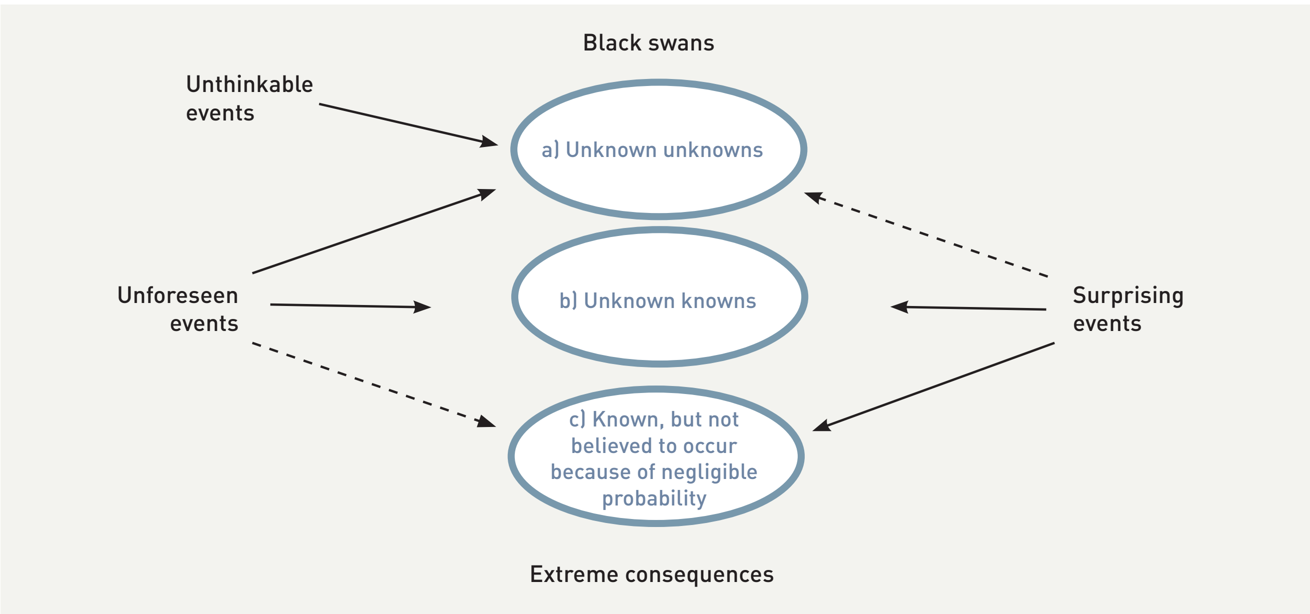

FIGURE 4: DIAGRAMMATIC PRESENTATION OF THE TERMS FOR DIFFERENT TYPES OF BLACK SWANS

then occurs actually comes as a surprise could be questionable. That explains the dotted lines in figure 4.

Similarly, what constitutes an unforeseen event is open to question. If we have identified an event but found its probability to be negligible, and this nevertheless happens, was it then foreseen? The answer is yes in the sense that its occurrence was predicted, but no in the sense that it was not considered probable.

Strictly speaking, categorising an unthinkable event as an unknown unknown would be reasonable. But arguments can again be produced for a different interpretation. Viewed from the perspective of the risk analysts, an incident in category b) can also be understood as unthinkable if a thorough job has been done to uncover all relevant events. The question is –unthinkable for whom. What ranks as surprising must always be understood in terms of the time and the person(s) we are speaking of. Figures 5-8 illustrate this.

Consider an activity (such as operation of an offshore installation) over a given future period – such as the following year. We let C represent the consequences of the activity in relation to the values which concern us (life and health, the environment, financial assets). What C will be remains unknown at time s; risk is present. See the discussion of risk in section 2.1. Suppose that a risk assessment of the activity is now carried out at time s. Time passes and C takes a value – normally without a major accident taking place.

But let us picture that such an event actually happens, as shown in figure 6. This is a result of a combination of events and conditions, and takes those involved in managing the activity by surprise. The accident is a black swan to them. But let us now adopt a macro-perspective by looking at a large number of such activities, such as the whole oil industry. Risk now relates to a major accident occurring somewhere or the other in this sector – neither where nor how is relevant.

A risk assessment is carried out. This may conclude that the probability for such an accident is relatively high. Consequently, an event of this kind cannot be described as a black swan if it actually occurs. See figures 7 and 8.

From a macro-perspective, a realistic analysis could indicate that we must expect a major accident to occur somewhere or the other over the next 10 years. However, no law says this will actually happen. One is not subject to a destiny (see the discussion on probabilities in section 2.2). The individual unit (organisation, company, installation) will work with the goal of preventing such an accident from actually happening. Good reasons exist for believing that this target can be met with systematic safety work. So such an event will normally come as a surprise, a black swan.

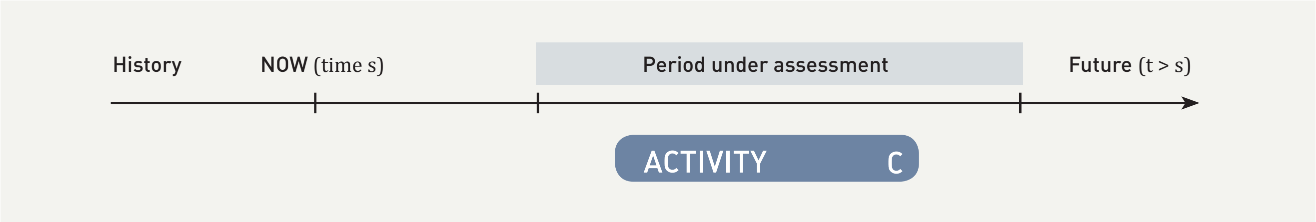

FIGURE 5: RISK IN RELATION TO THE TIME DIMENSION (C = CONSEQUENCE OF THE ACTIVITY)

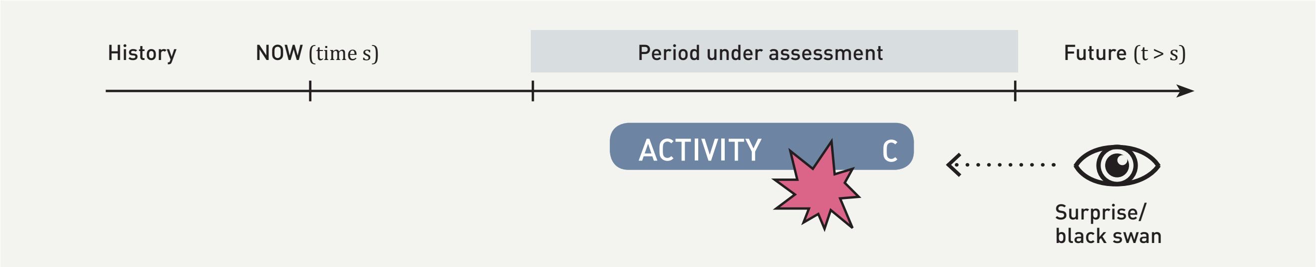

FIGURE 6: RELATIONSHIP BETWEEN RISK, BLACK SWAN AND THE TIME DIMENSION

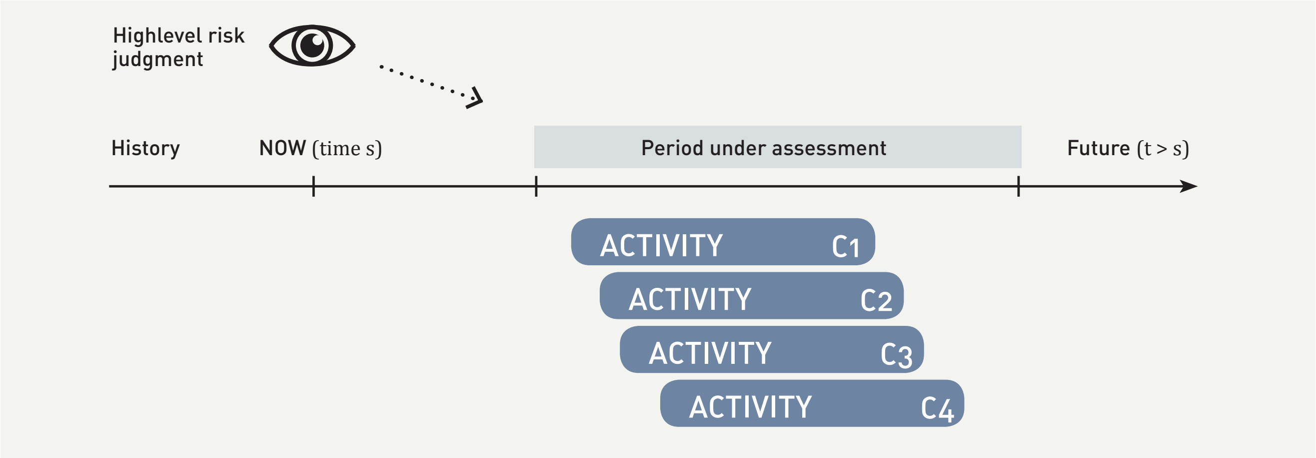

FIGURE 7: RISK IN RELATION TO THE TIME DIMENSIONS WHEN AN OVERALL PERSPECTIVE IS APPLIED TO A SET OF ACTIVITIES – THE WHOLE OIL INDUSTRY, FOR EXAMPLE

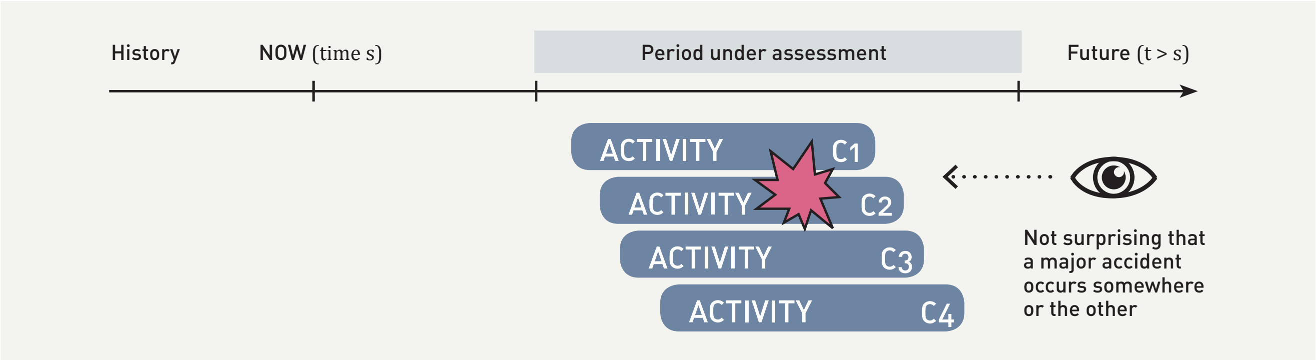

FIGURE 8: RELATIONSHIP BETWEEN RISK, BLACK SWAN AND THE TIME DIMENSION WHEN AN OVERALL PERSPECTIVE IS APPLIED TO A SET OF ACTIVITIES – THE WHOLE OIL INDUSTRY, FOR EXAMPLE

2.8 What do the five collective mindfulness principles mean?

THE FIVE PRINCIPLES

These are the five principles covered by collective mindfulness:

1. **preoccupation with failure**
2. **reluctance to simplify**
3. **sensitivity to operations**
4. **commitment to resilience**
5. **deference to expertise**

WHAT DO THESE MEAN?

The first principle means that we concentrate our attention on faults and weaknesses in the system. This is done to a great extent in the oil industry. Paying attention to failure signals is equally important. Reading these signals in the right way makes it possible to correct irregularities before anything serious happens.

This is the core of the principle concerning sensitivity to operations. When pressure in the process increases, we must understand what is happening and be able to make the necessary corrections. We must always take an integrated holistic view and be unwilling to simplify overmuch. Rules of thumb can be useful in some contexts, but must be used with caution because nuances exist in all circumstances and could be crucial for determining whether an event escalates. In one case, the risk could be assessed as acceptable simply by looking at its calculated probability. But this could be unfortunate if the knowledge it builds on is weak.

Undesirable events could happen regardless of how much work we devote to uncovering them and avoiding their occurrence. We must have barriers and emergency preparedness to meet them, both the known events and the unknown as far as possible. We must think resilience. A speaker giving a presentation must prepare for the possibility of difficult questions he/she has not thought about before. He/she must train for it. In the same way, one must think resilience in relation to the forerunners of black swans. They can occur. Being prepared for everything is naturally impossible, but thinking simply in terms of handling well-defined events will not be enough. We must also be able to cope with surprises.

This is what the fourth principle ¬– commitment to resilience – is about. The fifth – deference to experience – specifies that those with knowledge must be brought in when assessments are to be made and decisions taken. A speaker who does not know the answer to a question refers to others who do. Decisions in a process plant must be taken close to those familiar with the systems and who are knowledgeable. Sticking slavishly to a formal hierarchy in a crisis could be fateful when decisions are taken.

2.9 What are common- and special-cause variations? Why is it important to distinguish between them? What is the connection with risk?

The quality discipline places great emphasis on correct understanding, analysis and management of variation. To understand a system and improve it, we need to understand the reasons for variation in the process. The first ideas and analytical approaches in this direction came from Walter A Shewhart more than 80 year ago. He distinguished between two main types of variation – common cause and special cause (using the expression “assignable causes” for the latter, but special-cause variation is the usual term today). A system involving only common-cause variation is said to be under statistical control, and the outcome is controllable within certain limits. Special causes are identified with the aid of control charts, where limits are defined for when the process is/is not considered to be under control (usually specified as a standard deviation of +/- three from the average).

Common-cause variation is related to “aleatoric” (random) uncertainty in a risk context, and expressed with the aid of probability models. There is “epistemic” (knowledge-related) uncertainty about the common-cause variation, special-cause variation and results from the models in relation to the real variation (uncertainty about the model’s deviation).

In a risk context, the special causes relate to the surprises and the black swans. They call for immediate corrective measures. More fundamental change processes are required to do something about common-cause variations. Basic changes to technical solutions or the management structure may be needed. A special-cause variation could relate to a particular set of deliveries from a specific supplier or to a disloyal employee, and the measures needed to correct the process will generally be obvious. The distinction between a special-cause and a common-cause variation will always be open to question. An assessment must be made in each case. This is a matter of what strategies to adopt for the safety work and when fundamental changes are felt to be needed to boost the level of safety. One could choose, for example, to accept a substantial variation in the number of gas leaks in a process plant, and to respond only when the variation is extreme. An alternative would be to challenge the definition of a common variation at all times, and to seek to implement improvements in the underlying processes which generate this variation in order to achieve better results over time.

3 ANALYSIS AND MANAGEMENT

USE OF RISK MATRICES

USE OF SAFE JOB ANALYSES (SJA)

SURELY WE CANNOT MANAGE THE "UNTHINKABLE"?

HOW CAN KNOWLEDGE-BUILDING, EXPERIENCE TRANSFER AND LEARNING INFLUENCE RISK?

3.1 Use of risk matrices

Risk matrices have their clear limits in providing a picture of the risk associated with an activity. They cannot be used to draw conclusions about what is or is not an acceptable risk. For an assignment of the probability and consequence of an event, the strength of the knowledge on which it is build could be strong or weak. It would be an advantage to distinguish those events where the background knowledge is relatively weak, so that special care is taken in drawing conclusions based on the specification of such events.

MULTIPLE CONSEQUENCES

Furthermore, the description of consequences in risk matrices is often based on specifying one consequence per event, usually the expected consequence. However, a specific event will have a number of possible consequences with varying degrees of seriousness. This is picked up only to a limited extent by the typical consequence description presented in risk matrices. An example of an event with the potential for substantial deviance between expected and actual consequences is the "bus accident". The expected consequence of this incident could be a few injuries, but a large number of fatalities is also a conceivable outcome. The solution in such cases could be to split the undesirable event in two – "bus accident" and "bus accident with a number of fatalities", for example. This still represents a simplification, since the consequence of the accident could be one, two, three or more people injured, one, two or more people killed, or a combination of injured and killed. However, a balance will always have to be made between the level of detail and the amount of work needed.

As far as the strength of the background knowledge is concerned, the risk analysis could assume that everyone on the bus wears their seat belt – while the reality is that only a few do so when the accident occurs. The expected consequences does not reflect the failure to wear seat belts.

3.2 Use of safe job analysis (SJA)

The safe job analysis (SJA) is a well-established method for conducting risk assessment related to work operations. The method builds on risk descriptions based on events, consequences and probabilities. The same challenges noted in section 3.1, for example, are accordingly encountered here. SJAs can be improved by strengthening the risk description and thereby also the risk understanding. See the risk matrix example presented in section 3.1, where an assessment of the strength of knowledge is incorporated in the risk description. Greater attention should also be paid to aspects of surprise.

3.3 Is an enhanced risk perspective (mindset) significant for the way risk acceptance criteria are used?

Risk acceptance criteria are currently used in a fairly mechanistic way in the sense that calculated probabilities from risk analyses are compared with predefined acceptance limits for such probabilities. Conclusions about acceptable risk are then drawn on the basis of being above or below the acceptance limit.

This approach is characterised by i) a very arbitrary process and ii) providing the wrong focus. Where i) is concerned, the problem is that probabilities obtained from the risk analysis build on some background knowledge, and the strength of the latter is not reflected in the method. A probability can be calculated which lies below the acceptance limit, but the knowledge it builds on could be weak – and that must be reflected in the judgments. The following adjustments are proposed to the standard procedure.

If the risk is found to be acceptable by a wide margin in relation to the probability, this risk will be considered acceptable providing the knowledge is not weak (in such a case, a probabilistic approach should not be given much weight).

Should the risk be found acceptable in relation to the probability, and the knowledge is strong, the risk will be considered acceptable.

Where the risk is found to be acceptable in relation to the probability with moderate or narrow margins and the knowledge is not strong, it will be considered unacceptable and measures to reduce the risk will be required.

If the risk is found to be unacceptable in relation to the probability, it will be considered unacceptable and measures to reduce the risk will be required.

3.4 Are probabilities needed in risk management when talking about the surprising and the "unthinkable" (black swans)?

Probabilities are used to express variation and degrees of belief given the knowledge available. If our focus is black swans as defined in figure 4, we can quickly conclude that probabilities are not the appropriate tool for describing variation or likelihood. That is obvious where unknown unknowns are concerned. For category b), we lack knowledge about a type of events even if others possess it. Again, the probability concept is unsuitable. Where category c) is concerned, we have determined that the probability is so low that we disregard that it will happen. We must therefore look beyond this tool here as well.

To analyse black swans, we must go beyond the probability tool. We need other methods for analysing such events. In many respects, this is precisely what characterises these events – they are not picked up so well in the traditional approaches based on probabilities.

3.5 What actually are signals in relation to risk? Why is it important to understand these correctly?

Signals in relation to risk are an event, a change in condition and the like which indicate that something might happen. Consider a process plant where the pressure in a tank shows an abnormal rise, or where a valve starts to leak. We are all familiar with signals from our body which tell us we need to take a break from work. Failing to do so could easily prove serious. Paying attention to signals is clearly important because they allow corrections to be made and serious consequences to be avoided. Unfortunately, it is not always possible to read signals in the right way. This ability must be developed – a balance must be made between being alert and sensitive to signals and becoming oversensitive by reacting to all possible kinds of changes.

Achieving the right balance depends on a good risk understanding, and knowledge occupies a key place here. Some signs – signals – are not so important for safety that an immediate response is required, and good knowledge will reveal which these are. You can, for example, by all means push your body to make a further effort even if you feel tired at the start of a run. That poses no problem. But another type of symptom and signal can tell you that enough is enough, now is the time to stop. Most of us have experienced these, and know how to distinguish between them. Similarly, a good understanding of each system, experience and training will develop the ability to read signals correctly. This is a continuous process, which never ends. New and better insight can always be acquired. And a mindset is required which constantly seeks and encourages the acquisition of such insight.

3.6 What does viewing risk in relation to other considerations mean?

Why do we explore for oil, operate an offshore installation, invest in the stock market or step onto the Kjeragbolten?

Because we want to achieve something – find oil, make money and feel pride. Our main drive is not safety or low risk.

Safety and risk are aspects associated with the activities we pursue and, in certain cases, they can be too poor and too high respectively for us to initiate an activity. We are afraid that a major accident will happen, that we will lose money or fall off the rock and die. If the upside – the potential gain – increases, however, we could be willing to take or tolerate more risk.

We are all familiar with such mechanisms. When Norway's oil industry began, safety standards were not particularly high. Many people were injured or killed, but Norwegians accepted that as a society because the benefits for the country were so high. This was a matter of many jobs and huge government revenues.

If somebody offered you a million dollars you might have been ready and willing to step onto Kjeragbolten, even if you had originally refused. We cannot avoid the reality that an appropriate level of risk must also take account of the benefits. In practice, however, some limits always exist. Regardless of the size of the prize, few people in Norway are likely to play Russian roulette – which involves gambling with their own lives. The same holds true in industry, where there will always be a certain minimum level of risk and safety we regard as acceptable. However, it is impossible to define clear limits for this, since we must always view the level of risk and safety in a context which includes not only benefits but also ethical, cultural and political factors.

Photo: Kjeragbolten
About 1 000 metres above sea level at the innermost end of the Lysefjord near Stavanger, Kjeragbolten is a big boulder stuck fast in a deep and narrow crevasse. Posing on this rock has become a test of courage for many.

3.7 Surely we cannot manage the "unthinkable"?

That something is categorised as unthinkable and completely unpredictable can be used as an argument for doing nothing about it. Responsibility is repudiated on the grounds that the incident is anyway unthinkable before the event. But this is unacceptable. The first question, of course, is whether anything has been done to increase knowledge of the relevant phenomenon so that surprises can be avoided. That applies to all three categories – a) to c) – in figure 4. What is unknown to someone could be known to others, and what is unknown today could be known tomorrow. Knowledge changes all the time with the emergence of new data and information. If we cannot control the unthinkable, we can do a lot to influence and manage the associated risk – partly by increasing knowledge (converting the unthinkable to the conceivable) and by adopting measures to meet such events, like emphasising robustness and resilience.

Everyone will probably agree that managing the risk associated with drinking poison in the example discussed in section 2.6 is possible. A thorough risk analysis would pick up many types of incidents which might happen, but we quickly see that some events will always escape detection – either because we fail to see them or because they are identified but not given any weight as a result of negligible probability/risk.

The challenge is to identify which events we should concentrate our attention on, and how we can meet these and possible other events which might occur. This is risk management. Our way of thinking will help us to make the right choice here. We must take care not to think solely in terms of calculated probability, without taking account of the knowledge on which it builds.

Work is currently underway in a number of groups to improve risk analysis in order to take better account of the knowledge aspects and the potential surprises. No details will be provided here, but one example is that efforts are being made to improve the methods used to identify hazards. Creative methods exist to stimulate our thinking about how things might happen, but these are not much used in analyses today. Similarly, methods are being developed which can improve our risk understanding in specific circumstances, such as critical operations, on the basis of the risk mindset described here. Work is also underway to strengthen risk management in other ways. These include for instance evaluating and making proposals on improvements to various principles, methods and standards used in the industry, and further developing principles for precautionary thinking, continuous improvement and learning.

3.8 How are continuous improvement and anti-fragility an integrated part of the enhanced risk thinking ?

Requirements and compliance with these are important in safety work, but can easily act as a brake on improvement. Meeting the requirement becomes enough. Systematic processes for developing (tightening) the requirements have not been established. The ambition may be to rank as leading, outstanding or brilliant, but the leap from such ideals to practical everyday life can be a big one as long as continuous improvement is not an integrated part of the enterprise and risk management. Continuous improvement is a cornerstone of our thinking. This is linked to an understanding that variation, uncertainty and risk must be tolerated – even welcomed – to some extent if good (outstanding, brilliant) results are to be achieved. This is the main message in the latest book by the well-known Lebanese-American writer Nassim Taleb (author of the bestseller on black swans).

Taleb introduced a new term, “anti-fragile”, to describe how to live in a “black swan world” where serious unforeseen incidents occur. He makes the point that thinking protection and robustness is not enough – we must actually love some degree of “variation, uncertainty and risk”. Over time, this will lead to mastering and improved results.

That is exactly the matter. This is, after all, what you do when you exercise – you subject your body to stress and discomfort in order to improve your physical condition over time. To be a good speaker, you must expose yourself to risk, dare to break free, try new approaches and methods. You do not succeed every time, but know this is the right recipe for getting better.

That does not mean we should gamble with safety by failing to take compliance with requirements seriously. This is not the point. What it means is appreciating that stress and errors are not solely bad. We actually need a little of them. Who has not felt depressed at the start of a major mobilisation and change process which leads to substantial improvements? The stress and the errors must not be too great, otherwise things could go wrong, but a certain level of variation, uncertainty and risk is necessary to sharpen yourself and achieve the desired progress.

Many events which have led to near misses reveal poor risk understanding. New ideas and methods are developed to improve the understanding. Work on implementing these can cause stress and perhaps also lead to misunderstandings, lack of clarity and/or disagreements over the most important contributors to risk. Over time, however, these changes can provide a substantial strengthening in risk understanding and increased safety.

3.9 How will knowledge-building, experience transfer and learning influence risk?

In our thinking, knowledge is part of the description and understanding of risk. So knowledge-building, experience transfer and learning will clearly be key activities in risk management. While they are naturally also important in the traditional thinking, where risk relates to probabilities and history, the risk concept is unable there to reflect the knowledge dimension in the same way. An example of why we need a new risk mindset is provided in section 1.2 – the probabilities could be the same, but the knowledge they build on is strong or weak. The black swans of types b) and c) in figure 4 are of particular interest here:

b) unknown knowns – events which we have not included in our risk assessment, either because we are not aware of them or because we have not made a conscious assessment, but which are known to others.

c) known events considered to have a negligible probability, and which are accordingly thought not to happen.

One type of critical event could be known in certain parts of the company, but those directly involved lack this knowledge. That could reflect a failure in knowledge communication and experience transfer. Another event could be ignored because of its negligible probability even though more recent findings/research show that this is not appropriate. The knowledge is not good enough.

3.10 Why is theory important in relation to knowledge development and learning?

Theory is important for knowledge development and learning because it provides a method and structure for systematically comparing our observations with our understanding, and thereby for improving the latter next time round. This point is strongly emphasised in the quality discipline, where the four stages – plan, do, study, act – form an integrated aspect of improvement processes.

Meteorologists have theories and models which they use to be able to say something about the weather tomorrow and over the next few days. These are developed further by looking at what actually happens. If we see that certain measurements of the condition of a system are not moving in the desired direction, it could be appropriate to try to formulate a hypothesis about the cause and then use new measurements and experience to improve this hypothesis and thereby our understanding of the system.

Similarly, our risk thinking is built on a theoretical platform as described in this document and a number of scientific works. We believe this thinking represents a useful perspective and way of thinking about how to understand, assess and manage risk. The theory is substantiated by the arguments made, and can be adjusted and further developed through additional observations.

3.11 What is the problem with performance management? Can a focus on meeting requirements be reconciled with a continuous improvement process?

Many people, such as those in the quality discipline, are deeply sceptical about management by objective. This tradition emphasises the need to work on methods for improving processes rather than concentrating attention on setting numerical targets. The point emphasised is that a goal in itself provides nothing – it must be seen in relation to how it is to be attained. Such targets can easily lead to the wrong focus and self-deception. The important consideration becomes meeting the target, rather than taking an integrated and holistic view, and identifying what is actually best for the organisation as a whole. In a barrier context, it can be sufficient that all the performance requirements specified for the various barrier elements are satisfied. But this does not necessarily mean that the overall performance of the barrier system is good, that the overall barrier functions are ensured and that the risk is low. These overall dimensions are what matter, but they are harder to measure than the quality of the barrier elements. Attention can thereby easily become concentrated on the latter and on satisfying their requirements.

Herein lies a threat. Goals/requirements and satisfying/complying with these are given priority over the important processes needed to enhance understanding of the relevant system and to improve performance and results. Continuous improvement is much talked-about, but the actual emphasis will often be in practice on formulating and complying with requirements. The challenge is to secure systems which constantly encourage improvements. To succeed in this, the right balance must be found between thinking about detailed requirements and the overall picture.

Integrated and holistic thinking – the important things to develop so that the overall system becomes the best possible, reduces the overall risk and so forth – must be adopted for improvement processes to function efficiently. We can substantially sharpen requirements for a specific barrier element, but may obtain much more from an overall perspective through a moderate improvement to the element if this is viewed in connection with changes to other parts of the system. Integrated and holistic thinking is also important with a view to being able to deal with surprises and black swans. A concept like resilience must necessarily have holistic system perspective. The important consideration is not the way an individual barrier element works, but how the system functions and its ability to cope with the unforeseen.

4 EXAMPLES: BLACK SWANS AND ACTUAL INCIDENTS

WHAT CAN WE LEARN?

WHAT DOES THE PROJECT GIVE US?

4.1 What can we learn? What does the project give us?

SUMMARY

Let us recap a little from the previous chapters. A black swan is a surprising event (with extreme consequences) relative to our knowledge and beliefs. In other words, what ranks as a black swan will always depend on who experiences it and when. We recall how the black swan metaphor arose. Three hundred years ago, people in the western world believed that all swans were white. Then a Dutchman travels to Australia and discovers black swans. A surprise to us, but naturally not to the locals. The terrorist attack of 11 September 2001 was the same kind of event. This incident took most of us, including the authorities, by surprise. If the attack had not been unforeseen, something would have been done to stop it. But somebody clearly knew what would come. The event belongs to the unknown known category (type b) in figure 4.

FUKUSHIMA

Japan's Fukushima disaster on 11 March 2011 was also a black swan. The tsunami and its destructiveness took most people by surprise. But this belonged to a different category (type c) in figure 4. We knew that a tsunami of such a size could occur, it had taken place before. But its probability was judged to be so small that it would not happen. We cannot be prepared for everything, and must accept a certain level of risk. We can draw a parallel with Norway, which could suffer a tsunami from a submarine volcanic eruption in the Atlantic. However, this is not likely and Norwegians do not believe it to happen because they would otherwise have taken precautions. They are not doing so today. That is the case with many events. We ignore them because their probability is so small. It should be noted that the 11 September 2001 attack could also be categorised as type c). Before it happened, risk analyses had been carried out with scenarios like the one which actually occurred on that day. However, its probability was considered to be very low – people did not believe it would happen.

MACONDO

The Macondo accident with the Deepwater Horizon rig in 2010 again took many people by surprise. We experienced a set of events and conditions which, in advance of the disaster, must be said to have been extremely improbable from a risk analysis perspective. See section 2.6. Indeed, such an accident must take the operator by surprise or it would be impossible to talk about pursuing a prudent activity. But it nevertheless happened.

Similar type of surprises we experienced in relation to other extreme events. Take the 1988 Piper Alpha disaster in the UK North Sea, for example.

The most important events/conditions were:

- The control room operator started up a pump which was undergoing maintenance

- The fire water pump was in manual mode because of diving operations

- The neighbouring Tartan and Claymore facilities continued to produce to Piper Alpha

- The firewall was not designed to withstand the pressure created by a gas explosion.

All these things occurred. That would have been considered highly unlikely before the event.

SLEIPNER A

A further example is the Sleipner A concrete support structure which sank in the Gandsfjord outside Stavanger in 1991 during a trial submersion. The official explanation was erroneous calculation of forces as well as a reinforcement fault. This subjected the material to excessive stresses, causing it to fail and take on additional water. Once again, a set of events which was regarded before the event as highly unlikely – assuming they were even identified at all.

NEAR MISSES

The incidents described above were events which actually happened. But we have had many cases, of course, where the sequence of events came to a halt. The barriers have fortunately worked.

NEAR-MISS BLACK SWANS

In a schematic and simplified form, we can envisage that x events and conditions are required for an accident to occur, and that it will be averted if only y (less than x) of these happen. We get a "near-miss" black swan. This concept is understood to be a surprising event (viewed from the perspective of our knowledge and beliefs) which does not have extreme consequences but which could have produced them in only slightly different circumstances.

How can the black swan approach contribute in this context?

- It provides suitable concepts and a good grasp of what risk associated with black swans is about.

- It provides principles for analysis and management which can prevent (reduce the probability of) black swans, and facilitate the development of appropriate specific measures with such an effect.

Safety work involves identifying events and scenarios – based on hydrocarbon leaks, for example – and ensuring that barrier systems are in place to meet them if they occur. Many types of incidents occur over a year, but do not have serious consequences because the barriers work as intended. That was also the case with the near misses referred to above but, in a number of cases, the margins for avoiding an accident were small. Major accidents often happen because even more "surprising" events occur. Simply having an understanding of what this involves – with concepts such as surprising, unforeseen, risk and so forth – is important in itself for responding to these problems.

Knowledge and uncertainty are key concepts. Black swans are surprises in relation to somebody's knowledge and beliefs. In the 11 September example, somebody had the knowledge – others did not. Where as the Fukushima example is concerned, assessments and probabilities were the crucial consideration. But these are based, of course, on data and arguments/beliefs, so this is yet again about knowledge. Our thinking here must extend beyond today's practice and theory. We need new principles and methods. This project contributes with such principles, and presents examples of specific methods.

The project lays the basis for continued development of a set of specific methods and thereby for reducing the probability of black swans. We look more closely below at various subjects related to actual incidents, surprises and black swans.

4.2 Examples of how knowledge and lack of knowledge play a key role

A black swan as outlined above, is understood to be a surprising event (with extreme consequences) relative to our knowledge and beliefs. A number of examples have been provided above which show how knowledge and lack of knowledge is a key issue with regard to black swans. That also applies to near-miss black swans.

We can refer in this connection to the case of a hydrocarbon leak where the direct cause was the fracture of the bolts holding together a valve in a separator outlet. Sweating from the valve had exposed the bolts to produced water with a high chloride content and a temperature of roughly 120°C. That resulted in chloride stress corrosion cracking, weakening the bolts so that they ultimately fractured. Sweating had been observed externally on the valve. A risk assessment was then conducted, which concluded that the valve could be replaced during a later maintenance shutdown.

The people involved in that assessment and the subsequent decision were not aware of earlier experience with this type of corrosion and/or did not know that one assumption in the choice of material for the valve bolts was that they should not come into contact with produced water. However, this type of corrosion mechanism was well known at corporate level. Similar problems had occurred on the platform earlier, involving the fracture of two of four bolts because of the same corrosion mechanism. But routines had not been established to ensure that this knowledge was included in risk assessments related to sweating, and the experience was not applied when assessing the relevant case.

Another example is provided by the subsea gas blowout during a well operation on the Snorre A installation (PSA, 2005).

Integrity problems were identified in the well and taken into account during the original planning of the operation. The plan was amended and several risk assessments were cancelled. Combined with several other circumstances, these factors meant that the operation was conducted in a way which failed to take into account the known integrity problems.

The investigation report after the incident identified 28 nonconformities, including the failure of the well programme to reflect knowledge of the lack of integrity. In other words, some people knew about the integrity problems but not those involved in the final planning stage and the actual execution of the well operation.

This made a significant contribution to the blowout's occurrence. The Petroleum Safety Authority Norway (PSA) has described the incident as one of the most serious on the Norwegian continental shelf (NCS).

4.3 Examples of the surprising and unforeseen in incidents

Not all incidents come as a surprise. The perspective one hold is important for understanding this. That a blowout occurs somewhere or other once during the next decade cannot be characterised as surprising, but a specific accident or near miss on a given installation will normally take those responsible for this facility by surprise. The set of events and conditions is surprising. Macondo was mentioned in chapter 2.6. Another case in point is a hydrocarbon leak in 2012, which occurred in connection with testing two BOPs and discharged an estimated 3 500 kilograms. The process system is illustrated in figure 9.

Since different piping classes with varying design pressure are used on either side of the HCV and operating pressure exceeds design pressure from the HCV to the flare, ensuring that NC3 is not in the closed position when the other valves are open will be crucial. Such a design did not accord with more recent practice (NORSOK). The surprising and unforeseen aspect (for the personnel conducting the test) was thereby that operating the valves in the right sequence was critical. Nor was it clear to the personnel that NC3 was closed. This was assumed to be open, like HCV.

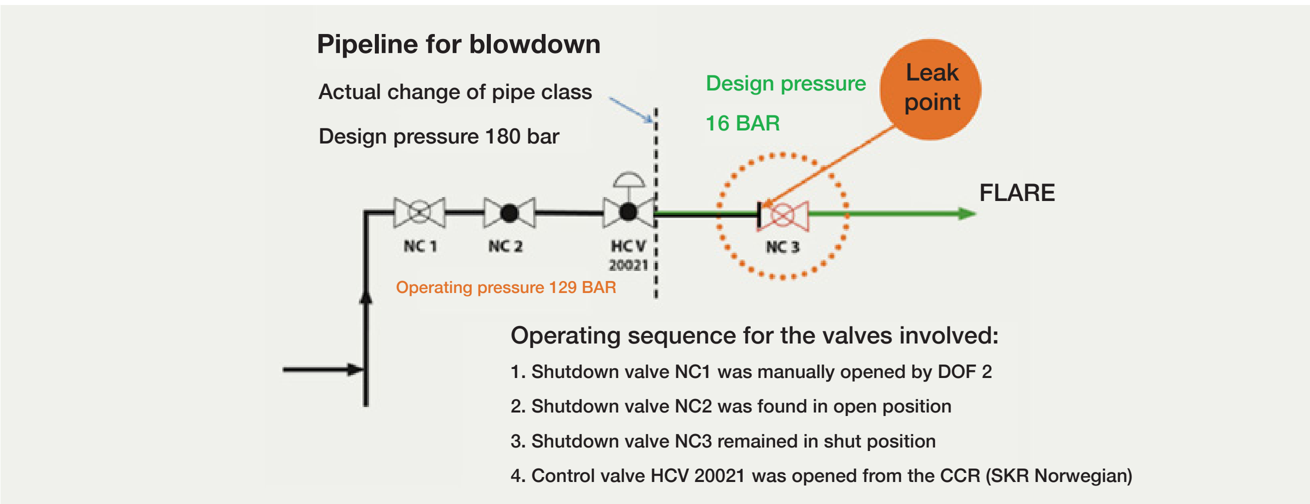

FIGURE 9: THE PROCESS SYSTEM

4.4 Examples of the mindfulness principles in view of actual incidents

A review of various serious accidents (the explosion/fire on Piper Alpha, the capsizing of Ocean Ranger, the Deepwater Horizon blowout and the Sleipner A sinking in the Gandsfjord) and near misses involving hydrocarbon leaks demonstrate up to several breaches of the mindfulness principles.

PREOCCUPATION WITH FAILURE (SIGNALS AND WARNINGS)

- Piper Alpha: insufficient attention was paid to a number of signals, such as the south-western flare making more noise than usual (Paté-Cornell, 1993).

- *Deepwater Horizon:* the report from the *Deepwater Horizon Study Group* (DHSG, 2011) points to basic failures in such areas as error and signal analysis. The organisation was also more concerned with success than with failure – it had "forgotten to be afraid".

- Hydrocarbon leak 1: the deviation of the design from more recent normal practice was not picked up.

- Hydrocarbon leak 2: sweating from the valve signaled that something was wrong, and a risk assessment was performed, but this failed to pick up the problem with the relevant type of corrosion.

RELUCTANCE TO SIMPLIFY

- Piper Alpha: both the two redundant pumps were maintained before the accident, but only things which were clearly broken got repaired. The rest does not appear to have been thoroughly checked (Paté-Cornell, 1993).

- *Ocean Ranger:* manual operation of the ballast control system failed because of inadequate competence/knowledge, which suggests that too much confidence was placed in the automated system (Vinnem, 2014).

- *Deepwater Horizon:* the report from the Deepwater Horizon Study Group (DHSG, 2011) in part describes simplified interpretations of results from critical negative pressure tests.

- Sleipner A: the contractor's advanced computer technology was supposed to be checked manually, but the routines failed (Stavanger Aftenblad, 2011).

SENSITIVITY TO OPERATIONS

- Piper Alpha: insufficient attention was paid to a number of signals, such as the south-western flare making more noise than usual (Paté-Cornell, 1993).

- *Deepwater Horizon:* the pressure rise in the drill string and a number of other signals were ignored for more than 40 minutes, and nobody noticed that a well kick was under way until drilling mud began to flow up onto the drill floor (DHSG, 2011).

COMMITMENT TO RESILIENCE

- Piper Alpha: the production team had minimum staffing when the accident occurred (Paté-Cornell, 1993).

- *Ocean Ranger:* the ballast control system was insufficiently robust, mainly because the single ballast control room was located inside an inner column of the rig. Nor was any reserve buoyancy built into the design as the final barrier against capsizing (Vinnem, 2014).

- *Deepwater Horizon:* the organisation was found to be far from robust – training and the support system for the operational team was poor, for example (DHSG, 2011).

- Hydrocarbon leak 2: the PSA noted a number of weaknesses in the way the incident was handled after its discovery, including lack of passive fire protection, lack of explosion resistance, failure to follow up identified non-conformities, and inadequacies with the emergency shutdown (ESD) system, maintenance of the process safety system, and the emergency preparedness and response plans.

DEFERENCE TO EXPERTISE

- *Ocean Ranger:* manual operation of the ballast control system failed because of inadequate competence/knowledge, which suggests that too much confidence was placed in the automated system (Vinnem, 2014).

- *Deepwater Horizon:* While initial well-design decisions are subject to rigorous peer review and subsequent amendments to a management of change process, typical changes in drilling procedures during the weeks and days before implementation are not subject to any similar processes. On Macondo, such decisions appear to have been taken by BP's drilling team on an ad hoc basis without any formal risk analysis or internal expert review (Graham et al, 2011).

- Hydrocarbon leak 1: the investigation report (PSA, 2012a) identifies significant weaknesses in expertise as well as in experience transfer and learning in the operations organisation after earlier incidents.

- Hydrocarbon leak 2: technical authority for materials was not included in the risk assessment underpinning the decision to postpone valve replacement. The reasons for this are not mentioned in the investigation report (PSA, 2013). In any event, it seems that the company's organisation possessed knowledge of the phenomenon but that those who had this were not included in the assessments which led to postponing valve replacement.

- The signal was detected and followed up, but ignorance of earlier experience with the relevant corrosion type among those conducting the risk assessment meant it was not interpreted to mean the leak was critical. In other words, the "signals and warnings" principle was not completely breached, but failure to involve the technical authority in the risk assessment represented a simultaneous breach of the "deference to expertise" principle.

4.5 Were risk analyses inadequate ahead of the actual incidents?

Most risk analyses begin by identifying events which could occur (events A), and their potential consequences (C) are then analysed and evaluated. Causes and explanations for the way such events could occur are also considered. In other words, traditional risk analyses ask the following questions:

- what could happen in the future, and what might it lead to?
- what could happen in the future, and what might cause it?

The future is uncertain in most cases, which means that we generally cannot know with any certainty what events that will actually happen and what their consequences will actually be. In such circumstances, risk analyses can be useful because their purpose is to systematise and describe knowledge of what the future might bring. Logical thought processes and concepts which describe uncertainty – such as probabilities and expected values – allow the analysis to systematise what we know, what we regard as probable and what assumptions support the reasoning. Describing knowledge of and beliefs about what might happen represents an important element in a risk analysis.

As mentioned above, a black swan is a surprising event with extreme consequences relative to the available knowledge and beliefs. Since the job of the risk analysis is to systematise knowledge, the natural question becomes whether a connection exists between black swans and near-miss black swans on the one hand and the lack of or inadequate risk analyses on the other. Are missing or inadequate analyses a characteristic of black swans and near-miss black swans?

To illustrate this, we will take a close look at which risk analyses have been carried out ahead of some relevant accidents/near misses.

- Loss of stability by Flotel Superior on 7 November 2012 on Norway's Njord field. In this incident, a loose anchor punctured the hull in eight places, water filled two tanks and the rig developed a significant list. Few formalised risk assessments had been carried out before the incident. The investigation report (PSA, 2012b) concluded that risk identification was lacking or inadequate.

- Oil spill on Statfjord A on 12 December 2007. A break in a loading hose allowed an estimated 4 400 cubic metres of crude oil to be pumped to the sea. This was the second-largest oil spill in the history of petroleum operations on the NCS. The investigation report (PSA, 2008) concluded that risk assessment of the whole loading system would have identified deficiencies in the technical solution.

- Hydrocarbon leak in 2012 when blowing down to the flare because pressure in a pipe segment exceeded its design level, reflecting a design weakness. According to the investigation report (PSA, 2012a), this weakness could have been discovered through a process Hazop or the like. It also concluded that carrying out planned risk assessments before the blowdown could have identified that the sequence for opening the valves posed a major accident risk.

All these incidents involved lack of or inadequate risk assessments. But it must be borne in mind that this conclusion reflects the wisdom of hindsight. Better risk assessments could clearly have identified incidents included in the near-miss and full black swan category, and led in a number of cases to changes and measures which could have prevented incidents from actually happening. But these analyses do not necessarily lead to the desired changes or measures. This is a question of risk management. After all, communicating and managing risk could involve major weaknesses. Moreover ¬– and this is important – it is always necessary to recognise that balances must be made, where different aspects are weighed against each other. We must accept a certain level of risk and we must think economics and efficiency. See section 3.7. When faced with a very large number of signals, it is not easy to judge which are the most important. However, history shows that poor risk understanding characterises many serious incidents, and better analyses and associated improvements in communicating risk can clearly make a positive contribution in many cases to ensuring that accidents actually do not to happen.

4.6 What was the state of risk understanding before the actual incidents?

The point has been made above that the identification of near-miss and full black swans depends on who and when – more specifically, the knowledge and beliefs possessed before the event. The terrorist attack of 11 September 2001 took most of us by surprise, but some individuals knew what was going to happen. In other words, an incident cannot be identified as a full or near-miss black swan without also clarifying whose perspective the assessment should be made from.

Similarly, a distinction must be drawn between different standpoints when assessing the state of risk understanding ahead of the incident. This cannot be described and assessed without clarifying whose risk understanding we are talking about. People can vary in their understanding of risk. The above-mentioned hydrocarbon leak 2 illustrates this. After sweating of produced water was observed, a risk assessment sought to evaluate whether postponing maintenance was prudent. The risk understanding of those who took the decision was that the leak of produced water did not have a major accident potential and that maintenance could thereby be postponed. However, the issue of chloride stress corrosion cracking was known to others in the organisation. So other people had a different, and better, risk understanding. But this did not form part of the basis when the decision was taken to postpone maintenance.

Another example is the 2012 hydrocarbon leak while blowing down to the flare. The design meant that the sequence for opening the valves towards the flare system was safety-critical. Based on the information in the PSA's investigation report, it is possible that the operators had an inadequate grasp of the criticality of ensuring that the final valve towards the flare was actually open, and thereby that the valves were opened in the right sequence. Since normal design practice had been changed in recent years, precisely to avoid this kind of criticality, others in the industry undoubtedly had a better risk understanding. Once again, we see that people can vary in their risk understanding and that inadequate understanding of risk was a contributory factor for the incident.

Piper Alpha is also a relevant example here. Before this accident, a consultancy warned the operator that a long-lasting high-pressure gas fire would have major consequences. At a subsequent meeting, however, the probability of this incident was considered to be so small that that it would not happen (Paté- Cornell, 1993). The Cullen commission (Cullen, 1990) concluded that the operator displayed a superficial attitude to assessing major accident risk. This suggests that the operator's risk understanding was inadequate. But it is important to note here that an actual accident does not necessarily happen because of inadequate risk understanding. Cases also exist of incidents where those involved unquestionably had or have a very good understanding of the risk, but nevertheless opted/opt to accept it.

A case in point is parachuting. Most of the people involved in this activity have made a detailed assessment of what could go wrong and what is required for that to turn into an accident. They have a good understanding of risk, but nevertheless choose to jump. This means that the risk is understood and accepted.

APPENDIX: CHECKLIST

A checklist of important aspects which need to be identified in order to take better account of the knowledge dimension and the unforeseen when carrying out risk analyses.

1. Has an overview been provided of the assumptions made? Related to system, data, models, expert assessments and so forth.

2. Has a risk assessment been conducted of deviations from the assumptions (individually and by looking at combinations of variations from several assumptions simultaneously)?

3. Have efforts been made to reduce the risk contributions from the assumptions with the highest risk of deviation?

4. Has the quality of the models used been assessed?

5. Have the model deviations (difference between the right value and the outcome of the models) been found acceptable?

6. Has an assessment been made of the strength of the knowledge which forms the basis for the probabilities?

7. Is this strength incorporated in the risk description?

8. Have efforts been made to strengthen knowledge where this is unsatisfactory?

9. Have special measures been adopted to identify unknown knowns – in other words, to learn about areas of expertise which the relevant analysis team lacks, but which can be found with others?

10. Have special measures been adopted to identify possible weaknesses (gaps) in the knowledge which the analysis team has based its analysis on?

11. Have special measures been adopted to assess the validity of assessments where events are in practice ignored because of their negligible probability?

12. Have people and expertise outside the analysis team been drawn on to identify such conditions as those mentioned above?

13. If expected values of a certain quantity have been specified, has the uncertainty related to this quantity been assessed (expressed by a 90 % uncertainty interval, for example, for this quantity)?

REFERENCES AND BACKGROUND LITERATURE

REFERENCES

Cullen, The Hon. Lord W. Douglas (1990). *The public inquiry into the Piper Alpha disaster.* London: H.M. Stationery Office. ISBN 0101113102. 488 pages, 2 volumes.

Deepwater Horizon Study Group (2011). *Final report on the investigation of the Macondo well blowout.*

Graham B, Reilly WK, Beinecke F, Boesch DF, Garcia TD, Murray CA (2011) *Deep water. The gulf oil disaster and the future of offshore drilling,* Report to the President. The National commission on the BP deepwater horizon oil spill and offshore drilling. Washington, DC.

Paté-Cornell ME (1993) *Learning from the Piper Alpha Accident: A Postmortem Analysis of Technical and Organizational Factors. Risk Analysis,* 13(2): pp 215-232.

PSA (2005). Gransking av gassutblåsning på Snorre A, brønn 34/7-P31 A, 28 November 2004. Stavanger: Petroleum Safety Authority Norway. http://www.ptil.no/nyheter/sammendrag-ptils-gransking-av-gassutblaasingen-paa-snorre-a-article1849-702.html

PSA (2008). *Oljeutslipp Statfjord OLS-A* 12.12.2007. Stavanger: Petroleum Safety Authority Norway.

PSA (2011). Ptils rapport etter Deepwater Horizon-ulykken: Vurderinger og anbefalinger for norsk petroleumsvirksomhet, 16/6-2011. http://www.ptil.no/nyheter/deepwater-horizon-ulykken-vurderinger-og-anbefalinger-for-norsk-petroleumsvirksomhet-article7889-24.html Stavanger: Petroleum Safety Authority Norway.

PSA (2012a). *Report from the investigation of a hydrocarbon leak on Heimdal, 26 May 2012.* Stavanger: Petroleum Safety Authority Norway.

PSA (2012b). Report from the investigation of the listing of Floatel Superior, 7 November 2012 on the Njord field. Stavanger: Petroleum Safety Authority Norway.

PSA (2013). *Hydrocarbon leak on Ula P platform 12 September 2012*. Stavanger: Petroleum Safety Authority Norway.

Sande T (2013, 1 October) *Misforstått risiko*. Stavanger Aftenblad, p. 20.

Stavanger Aftenblad (2011). http://www.aftenbladet.no/energi/olje/Derfor-sank-Sleipner-A-2853680.html

BACKGROUND LITERATURE

Aven, T (2014) *Risk, surprises and black swans: Fundamental concepts and principles of risk assessment and risk management.* Routledge Publishers, New York.

Aven, T (2013) *Practical implications of the new risk perspectives. Reliability Engineering & System Safety.*115, 136-145.

Aven, T., Røed, W. og Wiencke, H.S. (2017) *Risikoanalyse*, 2nd edition, Universitetsforlaget.

Aven, T. Krohn, B.S. (2014) *A new perspective on how to understand, assess and manage risk and the unforeseen. Reliability Engineering & System Safety*. Open access: http://www.sciencedirect.com/science/article/pii/S0951832013002159#

Aven, T (2015) *Implications of Black Swans to the Foundations and Practice of Risk Assessment and Management. Reliability Engineering & System Safety*, 134, 83–91. Open access: http://www.sciencedirect.com/science/article/pii/S0951832014002440

Flage, R & Aven, T (2009) *Expressing and communicating uncertainty in relation to quantitative risk analysis. Reliability & Risk Analysis: Theory & Application*, 2(2), 9-18.

Veland, H. og Aven, T. (2015) *Improving the risk assessments of critical operations to better reflect uncertainties and the unforeseen.* Safety Science, 79, 206-212.

Vinnem, J.E. (2014) *Offshore risk assessment : principles, modelling and applications of QRA studies.* 3rd ed. London: Springer.

Photos pp 8, 23, 35, 43: Tom Haga
Photo p. 33: travelinggold.com
Photo p 56: Statoil
English translation: Rolf E Gooderham
Design: Finnestad AS

Norsk olje&gass