SINTEF Industrial Management
Safety and Reliability

Address: NO-7465 Trondheim, NORWAY
Location: S P Andersens veg 5
Telephone: +47 73 59 27 56
Fax: +47 73 59 28 96

Enterprise No.: NO 948 007 029 MVA

SINTEF REPORT

TITLE

Organisational Accidents and Resilient Organisations: Five Perspectives

Revision 1

AUTHOR(S)

Ragnar Rosness, Geir Guttormsen, Trygve Steiro, Ranveig K. Tinmannsvik, Ivonne A. Herrera

CLIENT(S)

The Research Council of Norway

REPORT NO.	CLASSIFICATION	CLIENTS REF.	
STF38 A 04403	Unrestricted	Project no. 138459/230	
CLASS. THIS PAGE	ISBN	PROJECT NO.	NO. OF PAGES/APPENDICES
Unrestricted	82-14-02724-1	38 44 99	78
ELECTRONIC FILE CODE		PROJECT MANAGER (NAME, SIGN.)	CHECKED BY (NAME, SIGN.)
Perspectives Rev 1.1 Rapport.doc		Ragnar Rosness	Jan Hovden
FILE CODE	DATE	APPROVED BY (NAME, POSITION, SIGN.)	
	2004-01-15	Lars Bodsberg, Research Director	

ABSTRACT

Several major accidents are related to the interplay of organisational properties and technology. The aim of this report is to present a set of perspectives that can help us understand the organisational mechanisms related to major accidents. Five perspectives are discussed:

1. The energy and barrier perspective.
2. The theory of Normal Accidents.
3. The theory of High Reliability Organisations.
4. The information processing perspective.
5. A decision-making perspective.

The target groups of the report are researchers, students and advanced practitioners.

KEYWORDS	ENGLISH	NORWEGIAN
GROUP 1	Safety	Sikkerhet
GROUP 2	Organisation	Organisasjon
SELECTED BY AUTHOR	Accident	Ulykke
	Vulnerability	Sårbarhet
	Resilience	Robusthet

TABLE OF CONTENTS

Preface

This report gives an overview of theoretical perspectives on organisational accidents and resilient organisations. We hope that the overview will prove useful to researchers, students and advanced practitioners in search of a richer understanding of the mechanisms that make some organisations accident-prone, whereas other organisations experience remarkably few accidents.

In order to keep the task manageable with the available resources, we had to concentrate on theories related to major accidents. We hope it will be possible to extend the scope in a later version to cover more topics related to external threats and intentional damage. Even within this narrower scope, the overview is far from exhaustive.

This work was sponsored by The Research Council of Norway through the project *Risk and Uncertainty: management, understanding and adaptation*. More information on this project and other online publications can be found on our website www.risikoforsk.no. We want to thank Jan Hovden, Hilde K. Sæle, Terje Aven, Helene Blakstad and Camilla Knudsen Tveiten for constructive comments.

In Revision 1 of this report, the following changes have been made:

- The term ‘resilience’ is discussed in the Introduction.
- The term ‘safety functions’ has been replaced by ‘barrier functions’ in order to harmonise the terminology with an ongoing project on Health, Safety and Environment in the petroleum sector. We have also removed the proposal to limit the term ‘barriers’ to physical or technical measures.
- The fundamental concepts of Normal Accident theory are presented in more detail.
- New sections on “Mindfulness” and “Implications for risk reduction” have been added to Chapter 5.
- A section on “Emergency plans as fantasy documents” has been added to Chapter 6.
- Chapter 9 has been revised and slightly expanded.
- Sections 6.6 and 7.8 have been revised to make the relevance of the perspectives to the Åsta accident clearer.
- In addition, many minor changes have been made in order to clarify the presentation, remove spelling errors etc.

The changes do not affect the over all structure and approach of the report.

Trondheim, 2004-01-14

Ragnar Rosness

1 Introduction

1.1 Background

On December 25 1998 two workers were killed and eight others were injured in an explosion at an Esso gas plant at Longford in Australia (Hopkins, 2000b). Due to tight interconnections with two other gas plants, the gas supply to Melbourne was cut for two weeks. At the 30th of September, *The Age* brought the following witness from the operators that Esso blamed for the accident[1]:

> Things happened on that day that no one had seen at Longford before. A steel cylinder sprang a leak that let liquid hydrocarbon spill onto the ground. … Ice formed on pipework that normally was too hot to touch. Pumps that never stopped, ceased flowing and refused to start. Storage tank liquid that was normally stable plummeted …

The gas plant at Longford had enjoyed excellent LTI-rates for years prior to the explosion. However, the public investigation of the Longford explosion revealed several serious safety management deficiencies which contributed to the accident (Hopkins, 2000b). Many of the shortcomings could plausibly be viewed as results of extensive cost-cutting. This cost-cutting effort was actually part of a world-wide phenomenon, and thus paralleled the NORSOK initiative to reduce costs in the Norwegian petroleum industry (Hovden and Steiro, 2000).

The Longford accident was spectacular in its impact on regional gas supply, but it was not unique (Hopkins, 2000b). In fact, it makes sense to speak of a family of organisational accidents. These are often major accidents. They often come as "fundamental surprises" to many of the people that manage and operate the dangerous systems (Woods, 1990). However, several precursors are usually discovered if a public investigation is launched (Turner and Pidgeon, 1997). Some of the companies involved display excellent LTI-records. This suggests that the concepts and approaches that have been developed to handle minor accidents are not sufficient to understand and control the risk of organisational accidents.

We have recently witnessed major changes in technologies of hazardous systems, in the organisations that operate the systems, and in the political and economic environments of these organisations. The present dynamic society brings with it some dramatic changes of industrial risk management (Rasmussen and Svedung, 2000:10):

- A very fast pace of change of technology
- The scale of industrial installations is steadily increasing
- High degree of integration and coupling of systems
- A very aggressive and competitive environment.

[1] Cited from Hopkins (2000b:1).

Faced with these trends, many researchers and practitioners feel the need for new concepts that can help us understand how organisations become susceptible to organisational accidents. Many of us search for strategies and methods to build organisational resilience, i.e. to build organisations that are not prone to experience major accidents. It is important to understand why accidents occur in order to learn and benefit from them. But it is also important to study how the organisations handle their daily operations, correct deviations and learn from normal and abnormal situations.

The aim of this report is to present a set of perspectives that may help us understand the organisational mechanisms that may be involved in major accidents. We emphasise perspectives that have emerged from work on major accident risks in industry and transportation. Many challenges and issues at the organisational level are of a generic nature, i.e. they are common across many sectors. For instance, few if any organisations avoid the challenge of handling conflicting objectives. At the same time, the specific risk control strategies and measures need to be adapted to the threats and the constraints facing a specific organisation. For instance, civil aviation can "absorb" a few major accidents each year on world basis, whereas the nuclear power industry worldwide may need more than a decade to recover from a single event. This implies that civil aviation may afford to learn by hindsight to a greater extent than the nuclear power industry (Rasmussen, 1997). We suggest that the diversity is greatest and the scope for generalisation most restricted at the "sharp end" of the systems, i.e. physically and causally close to the actual hazards (Wagenaar et al., 1994).

1.2 Definitions and delimitations

At this point the reader may expect to find clear-cut definitions of "resilient organisations" and "organisational accidents". This leads to a bootstrap problem, since the perspectives we are going to present, focus on different aspects when they categorise accidents and account for organisational resilience.

As a starting point, we may define an accident as a sudden, unintended event or series of events where significant harm is inflicted on humans, the environment or material assets. This definition excludes intentional harm, such as terror, hacking or sabotage. The definition also excludes harm that occurs gradually, such as the long-term effects of continual emissions of toxic substances. The notion of "vulnerability" usually includes a system's susceptibility to intentional harm. Additional perspectives, or extension of the perspectives presented here, may be needed to adequately cover the issues related to an organisation's susceptibility to intentional harm.[2]

We may then consider Reason's (1997:1) conception of *organisational accidents*:

> [Organizational accidents] are the comparatively rare, but often catastrophic, events that occur within complex modern technologies such as nuclear power plants, commercial aviation, the petrochemical industry, chemical process plants, marine and rail transport, banks and stadiums.

[2] Readers interested in practical approaches to vulnerability and vulnerability analysis will find interesting material at the website http://www.ipk.ntnu.no/rams/ (partly in Norwegian).

> Organizational accidents have multiple causes involving many people operating at different levels of their respective companies. By contrast, individual accidents are ones in which a specific person or group is often both the agent and the victim of the accident. The consequences to the people concerned may be great, but their spread is limited. Organizational accidents, on the other hand, can have devastating effects on uninvolved populations, assets and the environment. … [Organizational] accidents are a product of … technological innovations which have radically altered the relationship between systems and their human elements.

Reason's definition is rather eclectic, and thus captures aspects that are important to several of the perspectives to be discussed.

Foster (1993: 36) defined *resilience* as an ability to accommodate change without catastrophic failure, or a capacity to absorb shocks gracefully. The word *resilience* conveys an ability to recover or spring back into shape or position after being pressed or stressed (elasticity), but also an ability to recover strength, spirits and good humour. In this report, the focus is on accidents. We can thus define a resilient organisation as *an organisation that has a capacity to accommodate failures and disturbances without producing serious accidents*.

Our main concern is how organisational accidents are related to the properties of the organisation during normal operations. We pay less attention to such issues as planning and training for emergencies.

The third limitation concerns the emphasis on the organisation's interactions with its environment. External threats have not been given much attention, although some general impacts of a competitive and dynamic environment are considered.

Since our topic is resilient *organisations*, we pay limited attention to issues related to the individual level, the regulatory level and the political level (see Figure 9, page 48). However, we acknowledge that organisations are parts of larger systems. For instance, vertical interactions between levels of decision making are discussed in Section 7.4. Moreover, problems and tensions at the organisational level are reflected at the level of individuals. Individuals may be caught in a double-bind situation where they face irreconcilable demands, or they face incomprehensible situations due to inadequate information handling at the organisational level.

Although this is a report on organisational accidents and resilient organisations, we have included a chapter on the energy and barrier perspective. The energy and barrier perspective pervades practical safety work, and therefore becomes a topic in some organisational theories of accidents. Moreover, the feasibility of controlling the risk by means of barriers may have an impact on an organisation's choice of risk control strategies (Rasmussen, 1994a).

We have not tried to be exhaustive, but rather to select a few complementary perspectives that are central in practical safety management or that have had a major impact on research and discussions in safety science.

1.3 Structure of the report

The main theoretical contributions to the understanding of organisational resilience are partly overlapping, partly complementary, and partly contradictory. We have not found any obvious way to systematise or synthesise this diversity. Trying to force everything into a single model would do injustice to the diversity and probably make for a complicated model that would be hard to communicate. Imposing a rigorous classification scheme on the theories is also problematic, since many theories are too complex and comprehensive to fit into a neat category.

Given this dilemma, we decided to group the material into five perspectives on organisational resilience. The perspectives represent different sets of assumptions and metaphors to make sense of organisations. The five perspectives are:

1. The *energy and barrier perspective*, according to which accidents can be understood and prevented by focussing on dangerous energies and means by which such energies can reliably be separated from vulnerable targets (Gibson, 1961; Haddon, 1970;1980). This perspective has been included because of its impact on practical safety management.

2. Perrow's theory of *Normal Accidents*, which explains some major accidents in terms of a mismatch between the properties of the technology to be controlled and the structure of the organisation responsible for controlling the technology (1984). This theory has provoked a lot of fruitful controversy, mainly because it concludes that some technologies should be abandoned in their current form because they cannot be adequately controlled by any conceivable organisation.

3. The theory of *High Reliability Organisations* (HRO) was developed partly as a reply to the challenge posed by Normal Accident theory (Rochin et al., 1987, LaPorte and Consolini, 1991). HRO theory is grounded in intensive studies of organisations that have demonstrated an outstanding capacity to handle fairly complex technologies without generating major accidents. Important concepts from this research tradition are *organisational redundancy* and a capacity of organisations to reconfigure in adaptation to peak demands and crisis.

4. *The information processing perspective*, taking Turner's theory of *Man-made disasters* as a starting point. (Turner, 1978; Turner and Pidgeon, 1997). In this perspective, an accident is viewed as a breakdown in the flow and interpretation of information that is linked to physical events.

5. A *decision-making perspective*, with a focus on the handling of conflicting objectives. Here we introduce Rasmussen's (1997) model of activities migrating toward the boundary of acceptable performance, as well as the notion of distributed decision-making.

The clustering of subjects we have done can of course be discussed, as can the way we have delimited the perspectives. We have tried to give a reasonably "rounded" presentation of each perspective. This implies that there are some overlaps, for instance between the information processing perspective and the HRO perspective.

Each perspective has been devoted a separate chapter. We have emphasised the following aspects of the theories in the presentations and the summary chapters (Chapters 8 and 9):

- Notion of immediate causation
- Notion of root causes
- Risk control strategies
- How can major accident risks be monitored?
- Critical assumptions
- How can organisational change influence risk levels?
- How can we learn from disasters and incidents?
- What is the relationship between minor and major accidents?

In organisational theory Morgan (1984) and Bolman and Deal (1986) have stressed the importance of combining perspectives in order to understand the organisations. A similar attitude is implicit in the way authors such as Reason (1997) and Hopkins (2000b) combine different perspectives in their discussions on organisational accidents. We think it is important to have a sense of all the perspectives to make better analyses and make decisions. We will not claim that one perspective is better than another is, but rather focus on what could be learned from the different perspectives.

In order to get a better grasp on the theories, we will examine how they can be applied to a recent train collision in Norway. We will summarise the event sequence and background of this accident in the following section.

2 The train collision at Åsta[3]

On 4 January 2000, a northbound train from Hamar was scheduled to stop at Rudstad station on the Røros line and wait for a southbound train from Trondheim to pass. However, the northbound train left Rudstad before the southbound train had passed. The two trains collided seven kilometres further north at Åsta station. The engine car of the northbound train was completely wrecked, while the steering car received minor damage and remained upright on the rails. The southbound locomotive train was severely damaged. The locomotive toppled over onto its side and the front carriage buckled and derailed. A major fire broke out immediately in the area around locomotive and the rest of the engine car. Few minutes later fire broke out in the front carriage and the fire eventually spread to the remaining two carriages. Out of a total of 86 people, 19 people were killed. The situation immediately prior to the collision is shown in Figure 1 on the next page.

The commission of inquiry identified two possible direct causes of the accident. They could not exclude the possibility that the exit signal for the northbound train at Rudstad station was green instead of red due to a short-term operational malfunction of the signalling and safety systems.[4] Neither could they exclude the possibility that the northbound train had driven out of Rudstad against a red exit signal.

[3] The summary of the accident is based on the report from the Commission of inquiry appointed by the Norwegian Government (NOU 2000:30).

[4] Several weak points were identified related to the safety and signalling system. It was known that the system could show an erroneous green signal for a moment due to the slowness of the relay system.

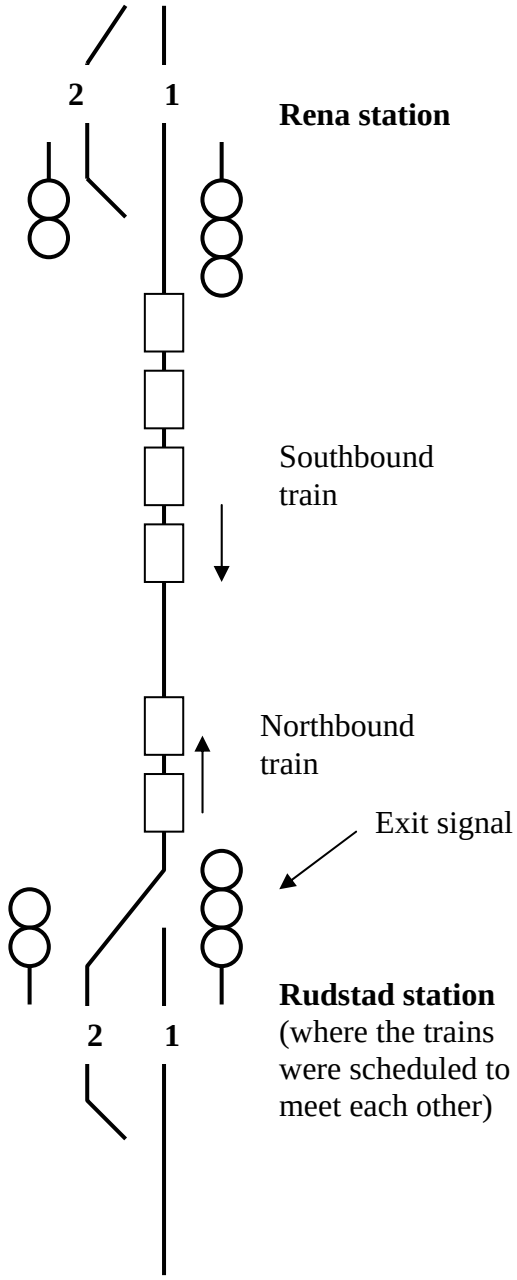

Figure 1. The situation immediately prior to the collision at Åsta. The northbound train had left from track 1 at Rudstad, forcing open the switch at the northern exit of the station. Not all signals are shown.

The Røros line did not have an operative Automatic Train Control[5] system (ATC) at the time of the Åsta accident. Automatic Train Control might have prevented the collision by automatically braking the northbound train if the exit signal was red when it started from Rudstad station. An audible alarm to warn of a train on collision course had not been installed at Hamar rail traffic control centre, from where the trains were directed. The rail traffic controller discovered a red

[5] The function of an Automatic Train Control system is to warn the driver and, if necessary, brake the train automatically if the driver exceeds speed limits or fails to brake the train adequately when approaching a red signal.

warning text on one of his displays about four minutes after the northbound train had left Rudstad and only *one* minute before the collision occurred. The Røros line is not electrified, so he could not prevent the collision by switching off the power supply to the trains. Train radios had not been installed on the Røros line. Both trains had mobile phones and both trains had reported in their phone numbers to the traffic controller at Hamar. The reporting of mobile telephone numbers was done to secure the traffic flow and service to the passengers. No safety grounds had been given for keeping a list of mobile telephone numbers, in spite of the requirements in the regulations about a rapid two-way contact between train and the control centre in case of an emergency. But the traffic controller on the previous shift did not add the numbers to the list. When the traffic controller on duty realised that a collision was imminent, he was not able to find the telephone numbers and contact the trains in time to prevent the collision.

In the following sections, we will present different perspectives on major accidents and discuss the Åsta accident on the basis of these perspectives.

3 Uncontrolled transfer of energy as the target of hazard control: The energy and barrier perspective

In cartoons, dangers are typically visualised in terms of energy – for instance a cliff and an abyss, a bomb or a bundle of dynamite with an ignited fuse. The idea that accidents can be conceptualised in terms of dangerous energies and inadequate barriers pervades theory and methods in the safety disciplines as well as practical safety work.

3.1 Energy transfer as the focus of accident research and prevention

Both practitioners and researchers are challenged by the diversity of accidents. The event sequences that lead to unintentional harm appears to be very different, the consequences range from trivial to catastrophic, and accidents occur in very different social and technological settings. Gibson (1961) introduced the energy model as a means to find some order in this perplexing diversity. He suggested that the most effective way of classifying sources of injury for research purposes is according to the forms of physical energy involved. The energy model was thus used by the medical discipline to systematise the analysis of accident causes in a way similar to that of analysing causes of diseases.

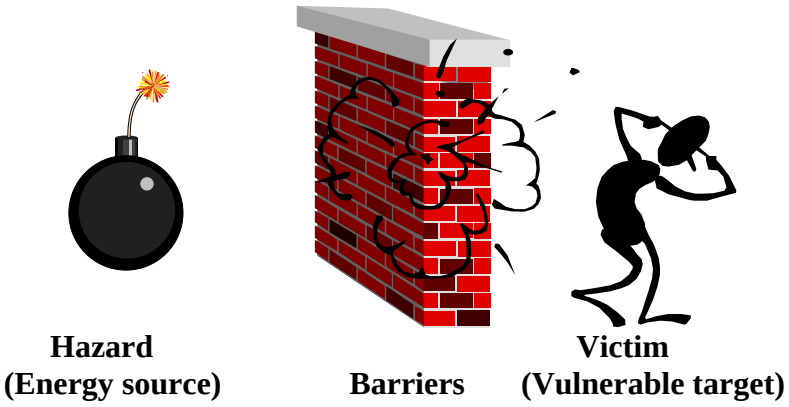

Figure 2. The energy and barrier model of accidents (adapted from Haddon, 1980).

William Haddon (1970, 1980) popularised the energy and barrier perspective and its implications for accident prevention. The basic idea is that accidents occur when objects are effected by harmful energy in the absence of effective barriers between energy source and the object (Figure 2). Haddon systematised known principles of accident prevention into 10 different strategies for loss reduction.[6] These are related to different points of intervention according to Figure 1:

1. Prevent build-up of energy (thermal, kinetic, or electrical); e.g. avoid car driving.

[6] Haddon also noted that each strategy could be turned into its opposite, i.e. a strategy to increase damage. In principle, the list may thus be used to identify strategies, which an actor may choose to intentionally inflict damage (e.g. sabotage). He even outlined how the principles could be applied to birth control.

2. Reduce the amount of energy; e.g. reduce the speed of vehicles.
3. Prevent uncontrolled release of energy; e.g. sanding and salting of roads.
4. Modify rate or distribution of the released energy; cars with shock absorbing zones, safety belts.
5. Separate in space or time, the victims from the energy being released; i.e. the use of sidewalks and the phasing of pedestrians and vehicular traffic.
6. Separate the victims from the energy by physical barriers; i.e. cars with safety cage.
7. Modify the qualities of the energy (the contact surface, subsurface, or basic structures), i.e. softening of hard objects in the car cabin.
8. Make the vulnerable target more resistant to damage from the energy flow; i.e. safety helmets.
9. Limit the development of damage; i.e. first aid.
10. Rehabilitate the victim(s).

Strategies 1, 2, 3, 4 and 7 are related to reduce *the hazard*, strategies 5 and 6 to *barriers*, while strategies 8, 9 and 10 are related to *protection and rehabilitation the victim(s)*. Higher-level loss control strategies may be formulated with reference to the ten basic strategies, e.g. "when feasible, prioritise risk-reducing measures directed at the hazard itself". Haddon argued that the larger the amount of energy involved, the earlier in the countermeasure sequence the strategy must lie.

3.2 What is a barrier? Barrier functions, barrier elements and barrier systems

The term 'barrier' has been given various definitions in the literature. In the basic energy and barrier model (Figure 2), a barrier is understood a means to separate a vulnerable target from a dangerous energy source. In concordance with this, Johnson (1980: 508) defined barriers as "The physical and procedural measures to direct energy in wanted channels and control unwanted releases." Kjellén et al. (1987) referred to the whole set of strategies proposed by Haddon as 'barriers'. These conceptions of 'barrier' thus include administrative measures such as procedures and work permit systems. However, some authors (e.g., Kjellén, 2000: 82) prefer to limit the term 'barrier' to *physical countermeasures* that intervene in the accident process to eliminate or reduce the harmful outcome.

The popular representation of the energy model (e.g. Figure 2) leads us to think of barriers as very concrete physical structures or devices. However, a *functional view* may be more productive when it comes to systematic loss control. A functional view implies that we think in terms of goals and means. We may think of a function as a *task* which is defined by one or more objectives to be achieved under specified conditions, for instance "prevent ignition of hydrocarbons after an uncontrolled hydrocarbon release in the process module"[7]. By taking a functional view, we thus focus on the *tasks that are necessary to adequately control a specific hazard.* These tasks may be performed by passive physical structures (e.g. fire proof walls), by active technical systems (e.g. the gas detection and emergency shutdown system on a production platform), or by humans, usually in interaction with technology and supported by procedures (e.g. the control of hot work so as to keep it separate from inflammable objects and substances). Thinking in terms of functions invites us to consider alternative means to implement a loss reduction strategy. For instance, if a gas detection system has to be inoperative during maintenance, an operator with a portable gas meter and radio communication with the control room operator may perform its task. Moreover, we need to consider that barriers can deteriorate and need to be monitored and maintained. This functional view fits well to the way Haddon formulated his loss reduction strategies, since each

[7] In this context, we should *not* think of a task as detailed, stepwise prescription of *how* a given objective is to be achieved. This would make us think in terms of part-whole-relations rather than goals-means-relations, and thus switch from a functional perspective to a system perspective (Rasmussen, 1986; 1997).

strategy is formulated as a task (e.g. 'Separate, in space or time, the victims from the energy being released'). The term 'safety function' is sometimes used in a sense similar to 'barrier function'.

Having defined a barrier function, we may identify the *barrier elements*, i.e. the hardware, humans, and software components (including procedures and routines) that are needed to implement the barrier function under given conditions. The totality of barrier elements that are necessary and sufficient to implement a given barrier function may be labelled a *barrier system*. The barrier system may thus be seen as the substratum or embodiment of the barrier function.[8]

Barrier systems are *open systems*. They do not function in isolation from their environment. Most technical devices can be disabled and need maintenance. Even passive barriers can fail due to human interventions. For instance, the performance of a fireproof wall may be dramatically reduced if it is penetrated by a cable bundle with flammable isolation. The distinction between physical and non-physical measures is thus not absolute. This is not just academic hair-splitting, because it points to a need for monitoring and maintaining barriers.

3.3 Defence in depth and organisational accidents

Haddon's model is relevant for the minor accidental event, as well as for major accidents. The prevention of accidents through *barrier functions* is an engineering approach and is a main principle behind safety in design (Kjellén, 2000:20). High hazard systems may employ several levels of defences in order to bring the total calculated risk to an acceptable level. For instance, a hydrocarbon processing plant (refinery or offshore installation) may include the following barrier functions related to hydrocarbon fires and explosions (Kjellén, 2000:85):

- Process control (automatic or manual);
- High quality containment;
- Gas detection and emergency shutdown;
- Isolation of ignition sources and ventilation;
- Fire detection and emergency shutdown;
- Area separation, fire/blast walls and passive fire protection;
- Active fire protection (e.g. deluge system);
- Provisions for escape and evacuation.

This hazard control strategy is commonly referred to as "defence in depth"[9]. A major accident in such a system is usually not caused by a single, isolated failure. This point is illustrated by Reason's (1997) "Swiss cheese model" (Figure 3).

[8] See E. Hollnagel: "Accidents and Barriers". Undated course note, Linköping: University of Linköping. Available at http://www.iav.ikp.liu.se/hfa/Coursemasters/Course%20materials/accidbarri.pdf

[9] The expression 'defence in depth' is of military origin. Lack of defences in depth can be demonstrated by the history of the Roman army. The Roman army was at most 300 000 soldiers deployed in the empire ranging from North Africa in south to northern Britain. The army was based on quick transport to areas where they were needed, but there was no real defence in depth. When the pressure towards the empire from the great people moving in Europe started, there were not enough soldiers to provide a defence in depth in several areas at the same time.

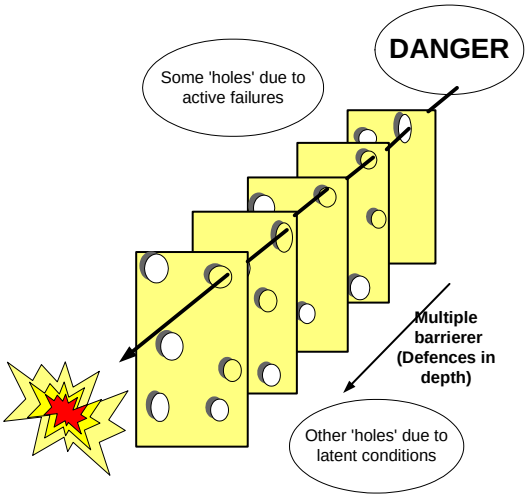

Figure 3. Organisational accidents involve the conjunction of failures in several defences. (Adapted from Reason, 1997:12)

Reason's (1997) figure shows an accident emerging due to holes in barriers and safeguards. In an ideal world all defensive layers should be intact allowing no penetration to happen. However, in the real world defences may deteriorate over time, such as the corroded sprinklers on the *Piper Alpha.* Modification or redesign may weaken or eliminate defences. Defences can be removed during calibration, maintenance and testing, or as a result of errors and violations. The control room operators of the Chernobyl nuclear reactor successive removed layers of defence in order to complete their task of testing a new voltage generator.

Reason distinguishes between *active failures* and *latent conditions*. Active failures are those that trigger unwanted events. They include violations and errors by pilots, doctors, and control room operators. These are the people in the operation or what Reason calls the "sharp end" of the system. Latent conditions do not immediately trigger an accident. However, they lie dormant in system and may contribute to a future accident. Examples are poor design, maintenance failures, poor and impossible procedures etc. These failures arise from top-level and strategic decisions and have indirect influence. Latent conditions can increase the likelihood of active failures.

In order to assess the effectiveness of defence in depth, it is not enough to assess the effectiveness of each barrier per se. We also need to consider *dependencies* among barriers. Dependencies will occur if two or more barriers can be weakened by the same event or condition. A failure in the electricity supply (blackout) may, for instance, leave several active technical barriers inoperative. This dependency may be increased if one or more backup energy supplies are likely to fail due to inadequate maintenance. Organisational conditions may thus create dependencies among barriers.

3.4 Analytical risk control

The energy perspective permeates major hazards control in the process industry and the nuclear power industry. For instance, the Quantitative Risk Analyses (QRAs) for Norwegian production platforms contain detailed models of the possible event sequences following hydrocarbon leaks in

the process area. These models emphasise the contribution of barriers and the number of persons that may be exposed to a fire or explosion. On the other hand, the conditions that lead to hydrocarbon leaks are not explicitly modelled.[10]

This is an example of *analytical risk control* (Rasmussen, 1997). Major explosions on production platforms or in nuclear power plants are rare and unacceptable events. Rather than learn from trial and error, we need to determine whether the risk level is acceptable before the system is built. The energy and barrier perspective allows us to do this, since we can design a system with several barriers and estimate the risk level based on assumptions concerning the effectiveness of these barriers.

In performing a QRA, one makes implicit and explicit assumptions related to the effectiveness of barriers. Assumptions concerning the effectiveness of barriers are thus implicit premises for a risk acceptance decision based on a QRA. The system can drift to a risk level not deemed acceptable if the effectiveness of barriers deteriorates significantly. The organisation should therefore establish programs to monitor and maintain barrier functions throughout system life. The methodology for continuous monitoring of the operational risk level is currently less developed than the methodology for design analysis (Øien, 2001).

An example of further use of the energy and barrier thinking is from the Swedish national road administration. The Swedish approach opens up for risk based setting of speed limits. Where there are a possible conflict between vehicle and pedestrian and no barrier is present, the speed limit should be maximum 30 km per hour. The fatality rate increases dramatically after the 30-km per hour limit is crossed in case of conflict. We know that unwanted energy can not always be avoided, so amount of energy should be reduced.

3.5 The energy and barrier perspective and the Åsta accident

We presented a brief description of the accident in Chapter 2. The commission of inquiry relied heavily on the barrier perspective in their analysis and evaluations. This is illustrated by the following excerpts from the investigation report (NOU 2000:30, p. 202):

> We know that technical systems can malfunction. We also know that people make mistakes. Consequently, there must be a safety system to ensure that individual faults do not result in accidents. ... Nonetheless, in the Røros incident, a signal failure or a mistaken observation by an engine driver led to a serious accident. ATC had not been installed on the Røros line … [Changes] were made without performing risk analyses for the individual change or for the Røros line. If the Norwegian National Railway Administration had done so, it would and should have been possible to see that an individual fault could lead to an accident. …
>
> With no barriers to prevent an emergency from arising, there should at least have been measures designed to avert it...

The investigation commission identified several defences which might have prevented the collision if they had been in place, for instance (NOU 2000: 30, pp. 171-175):

1. The presence of a train dispatcher at Rudstad would have reduced the risk that an error in the signaling system would develop to a critical situation. The driver of the northbound train would then have to obey signals given by the train dispatcher.

[10] In a production platform QRA, the frequency of hydrocarbon leaks is calculated by counting components that may leak and multiply by standardised leak frequencies.

2. A different departure procedure, which required the conductor to independently check the exit signal, might have allowed a red signal to be observed, or prevented a transient green flash from being mistaken for a steady green signal.

3. An Automatic Train Control system (ATC) would probably have stopped the northbound train in front of the exit signal, or at least a short distance behind the signal. The function of the ATC is to brake the train automatically in case the driver fails to observe a stop signal or a speed limitation.

4. An acoustic collision alarm in the train control centre at Hamar would have given the traffic controllers three to four minutes more time to notify the train drivers, provided that they reacted promptly to the alarm.

5. Rules as to how often the traffic controllers were to monitor their screens, combined with more staff at the control centre, could have reduced the time taken before an abnormal situation is noticed.

6. A train radio system would have allowed the traffic controller to reliably reach all trains on a railway section using a single number.

The accident commission concluded that the Røros line lacked adequate barriers against single failure accidents.

Here the problem of lacking barriers is demonstrated. Technical failures, deviations from procedure or erroneous actions happen. Therefore a barrier is needed. In the Norwegian rail, the Railway Inspectorate demands "the single fault principle". The principle states that no single technical failure, erroneous action or mistake should lead to fatalities or serious injuries.

3.6 Strengths and limitations of the energy perspective

Why has the energy perspective acquired a dominant position in major hazard control? There are several reasons:

- The energy perspective has proved very useful in hazard identification and as a basis for identifying hazard control strategies.
- The energy perspective is the basis for analytical risk control.
- It is possible to devise generic accident models by focusing on the uncontrolled release and transfer of energy.

Barrier functions allow individuals, groups and organisations *learning opportunities* which would not be available in an undefended system. An unintended shutdown in a process plant may be expensive, inconvenient and even hazardous, but it may force operators to practise skills that are not practised during normal operations. A shutdown may also provide information on the functioning of devices that are never used in normal operations, for instance emergency shutdown valves. This benefit of a well-defended system is often ignored, because it goes beyond the energy and barrier perspective.

The energy perspective may be most relevant for systems where the technical core and the hazard sources are well defined, physically confined and stable, for instance nuclear power plants or offshore oil production platforms (Rasmussen, 1994a). The scenarios following the release of a

major hazard in such systems are usually confined to one or a few paths (e.g. fire/explosion or structural collapse on an oil platform). In this case, quantitative risk analyses may emphasise the reliability of barrier functions.

In contrast, air transport is a distributed large-scale system. The functional technical core is divided among aircraft and infrastructures. Safe operations depend on the co-ordination of decentralised activities. Moreover, it is simply not feasible to design an aircraft strong enough to withstand a head-on mid-air collision. For these reasons, risk reduction efforts should emphasise preventing the release of hazardous energy, for instance by ensuring that critical systems are operative when they are needed.

Road transport could, however, benefit from more use of the energy/ barrier model. Haddon developed the energy barrier model for the road safety. And we have seen that the Swedish road authorities are using it more extensively.

Some authors include the notion of energy transfer in their definition of the term 'accident'.[11] From a physical point of view, any event has to involve a transfer of energy in order to be noticed. The occurrence of an energy transfer does not distinguish accidents from other events. Moreover, there are categories of accidents where the energy aspect is trivial. In an operating theatre, a pinprick with an infected syringe contains no more energy than a pinprick with a sterile syringe, although the former may cause a fatal infection (Hale, 2000). We may think of *information* – in the form of, e.g., DNA molecules, data viruses or computer bugs – as an alternative "medium" for the development of accidents. The barrier metaphor may prove useful even in the prevention of "information-driven" accidents, but this will require a different conception of barrier – one which is not linked to the energy model.

A limitation of the barrier model as it is used in QRAs of Norwegian production platforms is that factors influencing the initial event in analysis – e.g. hydrocarbon leaks – are not included in the model. This has an impact on what risk reducing measures are chosen if the QRA shows that the risk level is too high. The QRA does not give credit to measures devised to reduce the leak rate from a given component type. In this way, the energy and barrier perspective may lead to selection of sub-optimal measures for risk reduction in cases where new or improved barriers are not the most efficient measures.

Reason (Reason, 1997: 41) gave a picturesque example of soldiers that were *"killed by their armour"*. Heavily armoured French knights were thrown from their disabled horses by a storm of yard-long-steel-tipped arrows from English archers at Agincourt in 1415. The armour was so heavy, that they were unable to move or get on their feet in the mud. They were slaughtered by English foot soldiers equipped with mallets, spikes and daggers. In a similar manner, defences introduced to reduce the risk level may exacerbate an event under unfavourable conditions.

One common variety of this problem is the inflation of work procedures which can be observed in many organisations. Writing a new procedure is often perceived as a quick and inexpensive way to implement or reinforce a barrier function. One problem with this risk reduction strategy is that the total amount of procedures can become intractable. Operators may no longer find the time to identify all procedures applying to a given job, and the organisation may no longer find the resources to ensure that the total body of procedures is consistent, realistic and updated.

[11] For instance Johnson (1980: 507): Accident: An unwanted transfer of energy, because of lack of barriers and/or controls, producing injury to persons, property, or process, preceded by sequences of planning and operational errors, which failed to adjust to changes in physical or human factors and produced unsafe conditions and/or unsafe acts, arising out of the risk in an activity, and interrupting or degrading the activity.

Moreover, a very tight system of procedures may lead to more frequent conflicts between compliance with rules and efficient performance of the job. Such conflicts are often resolved through “silent deviations”. Routine violations of procedures become tacitly accepted practice. Discrepancies between rules and actual performance multiply, and activities may gradually drift out of control.

Active technical barriers may add to the complexity of the system, and thus increase the scope for maintenance-induced errors as well as operator errors. For instance, an automatic control system may be introduced to reduce a system’s vulnerability to operator errors. However, if the automatic control system fails, the operator may face an extremely difficult situation which he was not prepared for, since he no longer obtains the hands-on experience with the process (Bainbridge, 1987). Moreover, the automatic control system may add to the total complexity of the system, and thus make it more difficult to operate. Such paradoxes inspired Charles Perrow to formulate a theory of Normal Accidents, which we will discuss in the following chapter.

4 The challenge of interactive and tightly coupled technologies: Perrow's theory of Normal Accidents

Major accidents, such as the Three Mile Island accident, often come as fundamental surprise to the people that manage and operate the system (Woods, 1990). However, Charles Perrow (1984) insisted that some systems have structural properties that make such accidents virtually inevitable. He therefore labelled these fundamentally surprising events *"Normal Accidents"*. We will summarise his argument in this chapter.

4.1 Component failure accidents versus system accidents

Perrow (1984; see also Perrow, 1986:140ff) suggested that some major accidents are fundamentally different from minor events. Minor events are typically *component failure accidents*. They are caused by a failure of one or two components in a system, and they do not involve any unexpected interactions. The potential for component failure accidents can to a considerable extent be identified through standard risk analysis methods. For instance, in a Failure Mode and Effect Analysis (FMECA), the analyst considers one system component at a time, and identifies the possible failure modes. This analysis should capture a fair share of component failure accidents triggered by hardware failure, provided the analyst is able to cover each component, failure mode and relevant system state.

In contrast to component failure accidents, *system accidents* involve *the unanticipated interaction of several latent and active failures in a complex system*. Such accidents are difficult or impossible to anticipate. This is partly because of the combinatorial problem – the number of theoretically possible *combinations* of three or four component failures is far larger than the number of possible component failures. Moreover, some systems have properties that make it difficult or impossible to predict how failures may interact. We will return to these properties in the next section.

In the introduction to this report, we introduced Reasons's (1997) concept "organisational accident". What is the relationship between Reason's concept "organisational accident" and Perrow's concept "normal accident"? "Organisational accidents" are distinguished by the number of persons involved and the degree to which they belong to different parts of the organisation. "Normal accidents" are distinguished by the number of component failures involved and the quality of surprise – i.e. whether the event sequence was anticipated, or at least foreseeable. It seems plausible that these dimensions are correlated. An event sequence involving multiple failures seems more likely to involve several organisational units than a single failure accident. However, we should not jump to the conclusion that *all* "organisational accidents" are "normal accidents" and vice versa. The event sequence of some "organisational accidents" can, to a significant extent, be anticipated. For instance, the Åsta accident involved several agents belonging to several organisational units, and thus qualifies as an organisational accident. However, the event sequence did not constitute a fundamental surprise. It resembled the Tretten accident which occurred in 1975, and several persons had warned that this kind of accident might occur.

4.2 Complexity and coupling

Perrow proposed that some socio-technical systems have structural properties that are conducive to system accidents.

Some systems, such as major nuclear power plants, are characterised by *high interactive complexity*. These systems are difficult to control, not only because they consist of many components, but also because the interactions among components are *non-linear*. Linear interactions lead to predictable and comprehensible event sequences. In contrast, non-linear interactions lead to unexpected event sequences. Non-linear interactions are often related to feedback loops. A change in one component may thus escalate due to a positive feedback loop, it may be suppressed by a negative feedback loop, or it may even turn into its opposite by some combination of feedback loops. Such feedback loops may be introduced to increase efficiency (e.g. heat exchangers in a process plant). Even some safety systems may add to the interactive complexity of a system, for instance if overheating of a given component initiates automatic cooling. Interactive complexity makes abnormal states difficult to diagnose, because the conditions that cause them may be hidden by feedback controls designed to keep the system stable under normal operations. Moreover, the effects of possible control actions are difficult to predict, since positive or negative feedback loops may propagate or attenuate or even reverse the effect in an unforeseeable manner. *Unknown side effects* are another source of interactive complexity.

Another system characteristic that makes control difficult is *tight coupling*. Tightly coupled systems are characterised by the absence of “natural” buffers. A change in one component will lead to a rapid and strong change in related components. This implies that disturbances propagate rapidly throughout the system, and there is little opportunity for containing disturbances through improvisation. Tight couplings are sometimes accepted as the price for increased efficiency. For instance, Just-in-time production allows companies to cut inventory costs but makes them more vulnerable if a link in the production chain breaks down. In other cases, tight couplings may be the consequence of restrictions on space and weight. For instance, the technical systems have to be packed more tightly on an offshore platform than on a refinery, and this may make it more challenging to keep fires and explosions from propagating or escalating.

4.3 Organising for coupling and complexity

What we have presented thus far is a two-dimensional typology of socio-technical systems. Perrow used this typology to build an argument that some systems are intractable because they pose an organisational dilemma. The argument can be summarised as follows (see also Table 1):

1. *A system with high interactive complexity can only be effectively controlled by a decentralised organisation.* Highly interactive technologies generate many non-routine tasks. Such tasks are difficult to program or standardise. Therefore, the organisation has to give lower level personnel considerable discretion and encourage direct interaction among lower level personnel.

2. *A system with tight couplings can only be effectively controlled by a highly centralised organisation.* A quick and co-ordinated response is required if a disturbance propagates rapidly throughout the system. This requires centralisation. The means to centralise may, e.g., include programming and drilling of emergency responses. Moreover, a conflict between two activities can quickly develop into a disaster, so activities have to be strictly coordinated to avoid conflicts.

3. If follows from this that *an organisational dilemma arises if a system is characterised by high interactive complexity **and** tight couplings*. Systems with high interactive complexity can only be effectively controlled by a decentralised organisation, whereas tightly coupled systems can only be effectively controlled by a centralised organisation. Since an organisation cannot be both centralised and decentralised at the same time, systems with high interactive complexity and tight couplings cannot be effectively controlled, no matter how you organise. Your system will be prone to "Normal accidents".

Table 1. Organising for coupling and complexity.

Interactions Coupling	Linear	Complex
Tight	*Centralise to handle tight coupling!*	*Centralise to handle tight couplings AND decentralise to handle unexpected interactions!*
Loose	*Centralise or decentralise! (Both will work.)*	*Decentralise to handle unexpected interactions!*

Perrow applied his theory on the Three Mile Island accident. He concluded that the technology of the Three Mile Island power plant was so interactive and tightly coupled that it created the organisational dilemma described above.

4.4 Implications for risk reduction

According to Perrow, system accidents thus arise from a mismatch between the properties of a system (coupling and complexity) and the organisation controlling the system (centralisation versus decentralisation). The theory points to several risk control strategies:

1. With a complex system, you should try to reduce the degree of interactive complexity.

2. With a tightly coupled system, you should seek ways to loosen the couplings.

3. If you have to live with a high degree of interactive complexity, you should build a decentralised organisation.

4. If you have to live with tight couplings, you should centralise your organisation.

5. If your system has catastrophic potential, and you are not able to apply any of the above strategies, then you should discard your system.

Based on the last strategy, Perrow (1984) argued that some technologies, such as large, complex nuclear power stations and strategic nuclear weapon systems, should be discarded. The safety systems that are supposed to safeguard nuclear reactors create a degree of interactive complexity that may confuse operators and make system disturbances intractable.

4.5 A Normal Accident perspective on the Åsta accident

A few questions derived from Normal Accident theory may help us explore the Åsta accident in its structural context:

1. *Was the system characterised by a high degree of interactive complexity?* Traffic control on a railway system may be complex in the sense that there are many components such as trains, signals and switches. Moreover, the system is highly dynamic – its state changes from minute to minute. However, the interaction among components is largely linear. It is easy to predict what track a train will follow if you know the position of the switches. However, components within the technical systems that control the states of switches and signals may occasionally interact in a more complex manner. Because the electromagnetic relays do not always react instantaneously, a green light may occasionally occur for a second or two when the signal should have been red.

2. *Was the system characterised by tight couplings?* The basic idea of railways involves very tight physical couplings. Trains are confined to rails, and two trains on a collision course have no way to divert from their trajectories at the last moment. Moreover, it may be argued that the absence of effective communication equipment made the system more tightly coupled than it needed to be, since this reduced the scope for improvisation in an emergency.

3. *Was the organisation too centralised to cope with its interactive complexity?* We have argued that railway traffic control is mainly characterised by linear interactions. Linear interactions can, according to Perrow, be effectively controlled by centralised as well as decentralised organisations. The problem with electromagnetic relays is confined to a subsystem, and hardly calls for a decentralised organisation.

4. *Was the organisation too decentralised to cope with its tight couplings?* According to Normal Accident theory, a tightly coupled system such as a railway system can only be effectively controlled by a centralised organisation. However, it is not straightforward to judge exactly the degree of organisational centralisation in railway operations. Centralised Train Control (CTC) by definition implies a high degree of centralisation of traffic control decisions and of the operation of signals and switches. Very detailed rules and procedures also promote centralisation. Moreover, the train movements to a large extent pre-planned. However, once the northbound train had left Rustad station, it operated as an autonomous unit due to the lack of effective communication equipment. One may also question whether organisational fragmentation was a problem when NSB decided to change their departure procedure. Moreover, on may ask whether fragmentation of responsibility among decision levels may have contributed to the slow pace in introducing Automatic Train Control (ATC) on Norwegian Railways.

It is difficult to build a strong argument that the tragic event at Åsta was a clear-cut system accident in Perrow's sense of the word. One might argue that Åsta was a component failure accident, since a single active error was sufficient to trigger a catastrophe. It was known that trains occasionally pass a signal at danger. It was also known that this could lead to a catastrophe on a single-track railway without Automatic Train Control and without effective means for the Traffic Control Centre to detect and recover the error. The problem was not that the defences made the system opaque, but rather that the system lacked adequate defences. On the other hand, one might argue that the absence of effective means for communication between the traffic control centre and the train crew created a mismatch between the tightly coupled technology and an organisation where train crews temporarily operated as autonomous units.

A theory or perspective directs attention to some aspects of an accident, at the expense of others. The Normal Accident perspective does not focus on the absence of barriers designed into the system, such as Automatic Train Control in the case of the Åsta accident. Neither does it focus on the decision processes that led to a railway system that was highly vulnerable with regard to human error, although Perrow is highly aware that some systems remain error-inducing due to the distribution of power and interests among major stakeholders (Perrow, 1986:152f).

4.6 Strengths and limitations of Normal Accident theory

An important contribution of Normal Accident theory was to raise a discussion concerning the limits of safety in complex systems. Normal Accident theory thus inspired a research tradition on High Reliability Organisations, which will be discussed in the following chapter. The controversy following Perrow's book also inspired significant empirical research, for instance Scott D. Sagan's case study of the U.S. strategic nuclear weapons systems during the Cuba crisis (Sagan, 1993).

Perrow also drew attention to the possibility that some technologies may force us to adopt organisational structures and practices that are incompatible with central values in western democracies. In order to attain the degree of centralisation that is required to handle some tightly coupled systems, we may be forced to create work environments characterised by harsh discipline and very little autonomy.

Several objections have been raised against Normal Accident Theory:

- The notions of "interactive complexity" and "tight coupling" are so vague that it is difficult or impossible to subject the theory to empirical tests.
- It is difficult to derive a simple and effective prescription for assessing or monitoring major accident risk from Normal Accident theory, because it is difficult to measure or monitor such attributes as "interactive complexity" or "decentralisation".
- Analysis of recent major accidents suggests that most accidents result from other problems than a mismatch between complexity/coupling and degree of centralisation (Hopkins, 1999).
- Some critics find the suggestion that some technologies should be discarded too pessimistic, too fatalistic, or politically unacceptable.
- The assertion that an organisation cannot be centralised and decentralised at the same time sounds like a tautology. However, this assertion has been challenged by researchers that study so-called High Reliability Organisations (Weick, 1987). We will consider this challenge in the following section.

Some of Perrow's critics seem to assume that Normal Accident theory is only relevant to systems characterised by extreme interactive complexity and tight coupling. However, the theory has important implications for other organisations as well. For instance, Perrow claims that centralised control is necessary to handle a tightly coupled system. This implies that the operation of a major railway system, i.e. a tightly coupled technology, calls for centralised control. In practice, this may imply that the operational rules should be detailed, and not only specify functional requirements on task performance. This implication is neither obvious nor trivial, but it received some support in an interview study among operating personnel on Norwegian railways (Guttormsen et al., 2003).

Perrow's book was written at a time when technical systems were less integrated and competition was less fierce than it is today. We will probably be in a better position to appreciate the significance of his perspective after a few more years of exposure to current technologies and economic climate.

4.7 Further development of Normal Accident theory

Perrow seems to treat coupling and complexity as rather stable properties of sociotechnical systems. However, Weick (1990) argued that these attributes change during periods of crisis or high demand. For instance, the collision between two jumbo-jets at Tenerife airport in 1977 happened on a day when the airport was extremely crowded, it had to handle very large aircraft on a narrow runway, and visibility was poor. At least from the air traffic controllers' point of view, the system must have been more complex and more tightly coupled than on an ordinary day. Moreover, several errors occurred in communication between the tower and the two aircraft involved. Weick argued that these errors caused the system to become even more interactive and tightly coupled. In this way, Weick indicated that Normal Accident theory might be extended from a static, structural theory to a dynamic theory of how several failures can combine into an accident by making a system increasingly difficult to control.

Clarke and Perrow (1996) claimed that plans used to justify increasingly complex systems can impede organisational learning. We will return to this claim in Section 6.3, in the context of failures in information processing.

5 Organisational redundancy and spontaneous reconfiguration: The theory of High Reliability Organisations

5.1 "Working in practice but not in theory"

The previous chapter concluded in a pessimistic vein. Perrow claimed that highly interactive and tightly coupled technologies pose an intractable control problem. It is impossible to design an organisation, which is sufficiently decentralised to handle the interactive complexity, and at the same time sufficiently centralised to handle the tight coupling. This conclusion was challenged by a group of researchers who studied so-called High Reliability Organisations (HROs). Certain systems, such as aircraft carriers, nuclear submarines, air traffic control systems and nuclear power plants are only of benefit to society if they manage to deliver nearly failure-free performance (LaPorte and Consolini, 1991). At the same time, these organisations handle complex, demanding technologies and have to meet periods of very high peak demand.

The basic claim of this research tradition is conveyed in the title of a paper by LaPorte and Consolini (1991): *"Working in practice but not in theory"*. They claimed that some systems, which, according to Normal Accident theory, should be haunted by major accidents and not even able to produce anything useful, in fact do amazingly well. For instance, the conditions to be handled by the crew of an aircraft carrier was summarised as follows by a senior officer (Rochlin et al., 1987):

> So you want to understand an aircraft carrier? Well, just imagine that it's a busy day, and you shrink San Francisco Airport to only one short runway and one ramp and gate. Make planes take off and land at the same time, at half the present time interval, rock the runway from side to side, and require that everyone who leaves in the morning returns that same day. Make sure the equipment is so close to the edge of the envelope that it's fragile. Then turn off the radar to avoid detection, impose strict controls on radios, fuel the aircraft in place with their engines running, put an enemy in the air, and scatter live bombs and rockets around. Now wet the whole thing down with salt water and oil, and man it with 20-year-olds, half of whom have never seen an airplane close-up. Oh, and by the way, try not to kill anyone.

Given the inherent hazards, the complexity and tight couplings of this system, the safety records are remarkable according to HRO theorists. "For a deployment period of six months there will typically be over 10000 arrested landings with no accidents. Over 600 daily aircraft movements across potions of the deck are likely with a "Crunch rate"- i.e. the number of times two aircraft touch each other- of about 1 in 7000 moves". (LaPorte and Consolini, 1991: p. 21). The research challenge was not to explain why accidents occurred, but rather to explain *why so few serious accidents occurred.*

It is outside the scope of this report to evaluate the claims concerning the safety records of the HROs. We will concentrate on how HRO researchers explain excellent safety performance and on the relevance of this research for less exotic activities.

5.2 Organisational redundancy as a means to build fault tolerant organisations

Engineers are sometimes confronted with the task of building a reliable system from less reliable components. They achieve this by building in redundancy, i.e. but including extra (i.e. redundant) components that can take over in case a critical fails. Thus the braking system of a car comprises two separate hydraulic circuits, although a single circuit could do the job perfectly well. LaPorte

and Consolini (1991) found that the HROs used the principle of redundancy to derive highly reliable performance from less than perfect human beings. The organisation had its share of errors and deviations, but unlike less reliable organisations, it was able to correct the errors immediately. Crew members with overlapping tasks and competence. They had eye-to-eye contact and could easily communicate with each other. They were thus able to spot each other's slips and mistakes, and the culture supported intervention to recover the errors. For these reasons, nearly all critical errors were recovered.

Rosness et al. (2000) termed this error recovery capability "organisational redundancy".[12] They proposed that organisational redundancy depends on (1) structural/instrumental preconditions and (2) cultural preconditions (Figure 4). The *structural/instrumental dimension* of organisational redundancy concerns the personnel's possibility of direct observation of each other's work, overlapping competence, and overlapping tasks or responsibilities. Roberts (1989) and Bierly and Spender (1995) noted that HROs devote much attention to the development and maintenance of individual and collective competence. Some organisations build structural robustness by distributing veto powers, particularly in situations where inaction is a safer state than action (Schulman, 1993). Another important aspect of this dimension is the diversity and quality of communication channels. Weick (1987) argued that rich communication, for instance face-to-face discussion, is in general more powerful in promoting reliability in a complex system than sparse communication such as formal written messages.[13]

The *cultural dimension* of organisational redundancy concerns the capability and willingness to exchange information, provide feedback, reconsider decisions made by oneself and colleagues, and intervene to recover erroneous actions. LaPorte and Consolini (1991: 29) observed apparently contradictory production-enhancing and error-reducing activities in HROs. People reported errors without encouraging a lax attitude toward the commission of errors. They took initiatives to identify and improve flaws in Standard Operating Procedures. Error avoidance was achieved without stifling initiative or operator rigidity. People monitored each other's performance without counterproductive loss of operator confidence, autonomy and trust. In critical situations the crew gave each others orders and instructions independent of the military rank.

[12] Rosness et al. proposed the following definition of organisational redundancy: "By 'organisational redundancy' we refer to *co-operation patterns that allow the organisation as a whole to perform more reliably than each individual operator.*"

[13] There are, however, some situations where restrictions and standardisation may be necessary to prevent critical misinterpretation, e.g. in communication between pilots and Air Traffic Control operators.

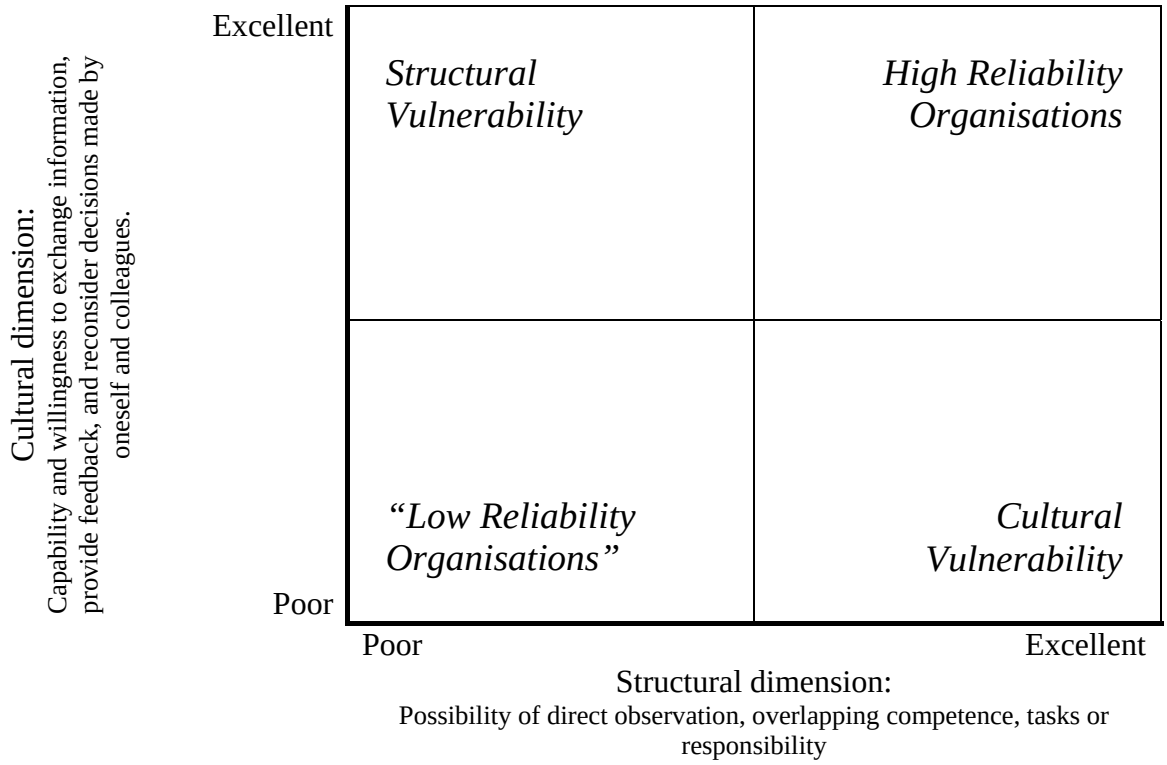

Figure 4. The two dimensions of organisational redundancy.

5.3 Spontaneous reconfiguration of the organisation

LaPorte and Consolini (1991) demonstrated another important aspect of HROs, which Perrow had not paid attention to. HROs are able to reconfigure spontaneously in during demanding operating situations and crisis. The aircraft carrier had a traditionally military system with commando lines clearly defined. But in situations with peak demand, the HRO changed into a more flexible and resilient pattern. Informal authority was granted on the basis of competence rather than rank. Interaction style became informal. In an air traffic control centre, controllers may even change the distribution of tasks and their physical location in the control room. For instance, extra personnel may join operators with particularly demanding tasks in order to provide "an extra pair of eyes".

Situational factors are well known in organisational theory. Burns and Stalker (1961) demonstrated that a hierarchy could be efficient during stable and predictable conditions, but inefficient under dynamic and changing condition. Bolman and Deal (1986) provide an example of a commando group during World War 2. Needless to say, these operations carried out behind enemy line were risky. On the other hand, these groups tried to avoid direct confrontation and therefore minimise the interactions with the enemy. The most successful commando unit during the allied operations, was very participate and involving all the members when planning the mission. The best suggestion was sought after. However, during the operations the commanding officer was the one handling all the decisions and had the right and duty to improvise when needed. One of the main points of Bolman and Deal (1986) is that organisations must be understood and managed by the combination of different perspectives and frameworks.

5.4 Culture as a means to build organisations that are both centralised and decentralised

A central claim in Perrow's theory of normal accidents is that an organisation cannot be centralised and decentralised at the same time. This is the reason why interactive and tightly coupled organisations create an irresolvable dilemma, see Section 4.3, p. 24. Weick (1987) challenged this claim. He argued that culture can impose a high degree of order and predictability in an organisation, and thus substitute formal means of centralisation such as a tight control structure or detailed operating procedures. Weick suggested that story-telling is one of the most important means for culture co-ordination, since natural language is a far richer means of communication than formal procedures, accident report forms, or statistics.

5.5 The notion of "mindfulness"

Weick and Sutcliffe (2001) introduced the notion of "mindfulness" to capture prominent characteristics of HROs. They state that HROs accept the fact of failures, that there is no perfection of zero errors. If errors are inevitable, the organisation needs to develop skills to detect errors and to contain these errors at early stages. Weick and Sutcliffe also emphasise the willingness of people to revise their own expectations. The idea of mindfulness is the equivalent of continuous surveillance of the existing situation based on expectations, updating these expectations based on new experiences and willingness and capability to invent new expectations. Mindfulness concerns detection and containment of unexpected events that could appear everywhere in the organisation.

Weick and Sutcliffe identified five elements that characterised mindfulness. These elements are organised under two tittles: *Anticipation and awareness of the unexpected* and *Contain the unexpected,* see Table 2.

Table 2. Elements of "Mindfulness". Summarised from Weick and Sutcliffe (2001).

Anticipation and awareness of the unexpected	**Description**
Preoccupation with failure	People in HROs know that all potential failures modes have not been experienced or exhaustively deduced. Because the cost of the failure is so high, people in HROs look for symptoms and encourage reporting of errors.
Reluctance to simplify interpretations	Simplify less and see more. Simplifications could produce blind spots, HROs use people that represent different functional background to expand the organisation's sensing mechanisms.
Sensitivity to operations	Normal operations can reveal deficiencies – free lessons could be learned. This allows early problem detection before problems become too substantial.
Contain the unexpected	**Description**
Commitment to resilience	HROs are not error free, but errors do not disable the system. People in HROs with varied experience come together as the situation demands, it increases the knowledge and actions can be brought to solve the problem
Defence to expertise	Decisions are made in the front line. Decisions migrate to the persons with experience and expertise to solve the problem.

Weick and Sutcliffe present "mindfulness" as a more or less universal cluster of features characterising HROs. They give little attention to the possibility that the means to achieve high

reliability may depend on factors such as the properties of their technology, as suggested by Normal Accident Theory.

5.6 Implications for risk reduction

Weick and Sutcliffe (2001) proposed a broad set of practices to develop mindfulness. The following are examples of practices that aim to *enhance awareness and anticipation to detect the unforeseen:*

- Leaders help employees to cope with conflicts, preserving a balance of values.
- Restate goals in the form of things that should not happen. This will provide more focus on the unexpected, disconfirming expectations and on issues of reliability.
- Remember that mindfulness takes effort. It is difficult to pay more attention to failures than to success. Look at failures, assume nothing, look closely at the work involved in the problem, brainstorm a resilient response, and pinpoint the expert in handling the problem rather the person accountable for the problem.
- Create awareness of vulnerability. What is risky around here? People need to worry about vulnerability as this increases opportunities for learning, commitment to reliability and accepting the fact at even though the system is understood, it can fail.
- Present mistakes as opportunities for enlarged learning and deeper understanding.
- Create an “error friendly” learning culture, by promoting behaviours such as seeking feedback, sharing information, asking for help and talking about errors.
- Encourage alternative frames of reference, strengthen fantasy.
- Communicate, promote scepticism, seek out bad news, test your expectations.
- Welcome uncertainty, treat all unexpected events as information and share them.

A second set of practices concerns *containment once the problem is evident:*

- Ambivalence builds resilience. Begin to contain the event by doing what experience tells you but watch for what you have not seen before and deal with it immediately.
- Use rich communication media, e.g. face to face communication rather than e-mails. Make sure that everyone’s voice is heard.
- Think out publicly when you question your categories, spot limitations and see new features of the context.
- Enlarge competencies and responses repertoires, then people will be able to see more hazards.
- Create flexible decision structures, let the problem migrate to the people who have the most expertise to deal with the problem.
- Accelerate feedback so the initial effect of the attempted improvisations can be detected quickly and the action altered or abandoned if the effects are making things worse
- Balance centralisation with decentralization, maintain local and centralized capacity for detection of problems this will enhance awareness in the organisation.

A more comprehensive description of these practices is found in Weick and Sutcliffe (2001:159ff).

5.7 HRO theory and the Åsta accident

But how can the Åsta accident be seen in an HRO theory? One of the issues that were discussed following the Åsta accident was whether the departure procedure for trains should formalise double-checking of the exit signal at stations. Until 1997, the departure procedure stated that the train driver and the main conductor has to check independently that the train has received a green exit signal before it starts after a stop on a station. This was an attempt to build organisational

redundancy. From 1997, the procedure was changed, so that this responsibility was only assigned to the train driver. This change was implemented by the Norwegian National Railway Administration and by the Norwegian State Railways even though the Norwegian Railway Inspection Authority refused to accept the change. In spring 2001, the Norwegian National Railway Administration decided to revert to a departure procedure that involves double-checking of the exit signal.

As an aside, we may note that before Centralised Train Control system was installed on the Røros line, traffic control was based on an old-fashioned, manual system. The railway section between two stations was reserved for a specific train by an exchange of telegraph signals between the two stations. When the train had passed the section, a new exchange of telegraph signal served to release the section, so that it could be reserved for another train. A detailed human reliability analysis of this apparently antiquated system has shown that extensive organisational redundancy was built into the procedures.[14] For instance, the train drivers would know in advance on what station they had to wait for a meeting train. They would thus refuse to start the train if the station personnel erroneously gave them permission to leave a station before the meeting train had arrived. This illustrates that organisational redundancy is not restricted to high-tech systems. Moreover, it shows that organisational redundancy can in some cases be highly formalised. The success of this system is attested by the fact that catastrophic train collisions were very rare events, even when fallible humans had to carry out tasks that are now performed by failsafe interlock systems.

5.8 Normal Accident theory versus High Reliability Organisations

HRO theory emerged partly as a response to Perrow's pessimistic view on the feasibility of reliably operating highly interactive and tightly coupled technologies. Sagan (1993) summarised the contrasting position as shown in Table 3.

[14] This claim is based on an Action Error Mode Analysis performed by SINTEF and the Norwegian National Railway Administration.

Table 3. Competing perspectives on safety with hazardous technologies. Adapted from Sagan (1993:46).

High Reliability Theory	Normal Accidents theory
Accidents can be prevented through good organisational design and management.	Accidents are inevitable in complex and tightly coupled systems.
Safety is the priority organizational objective.	Safety is one of a number of competing objectives.
Redundancy enhances safety: Duplication and overlap can make " reliable system out of unreliable parts."	Redundancy often causes accidents: it increases interactive complexity and encourages risk-taking.
Decentralised decision-making is needed to permit prompt and flexible field-level responses to surprises.	Organizational contradiction: decentralization is needed for complexity, but centralization is needed for tightly coupled systems.
A "culture of reliability" will enhance safety by encouraging uniform and appropriate responses by field-level operators.	A military model of intense discipline, socialisation, and isolation is incompatible with democratic values.
Continuous operations, training, and simulations can create and maintain high reliability organizations.	Organizations cannot train for unimagined, highly dangerous, or politically unpalatable operations.
Trial and error learning from accidents can be effective, and can be supplemented by anticipation and simulations.	Denial of responsibility, faulty reporting, and reconstruction of history cripples learning efforts.

In an attempt to validate the conflicting claims of Normal Accident theory and HRO theory, Sagan (1993) examined the operations of the U.S. nuclear forces during the Cuba crisis. This case was selected because the public safety records of the U.S. nuclear forces were excellent, the Defence Department claimed that risk accidentally releasing a nuclear attack was virtually zero. The idea was to submit Normal Accident theory to the "tough test". Through his investigations, Sagan discovered several serious incidents. These incidents revealed the types of problems that Perrow claimed would haunt highly interactive and complex systems (see the right column in Table 3). He found that extreme discipline could "encourage excessive loyalty and secrecy, disdain for outside expertise, and in some cases cover-ups of safety problems, in order to protect the reputation of the institution" (p. 254). Although the official commitment to avoiding mistakes, miscalculations or misunderstanding was very clear, lower level decisions repeatedly reflected other priorities, such as the maximisation of military preparedness. Sagan was particularly concerned that organisational learning can be constrained by the strong disincentives against exposing serious failures.

Rasmussen (1994a) argued that Normal Accident theory and HRO theory may be more compatible than Sagan's analysis suggests. Perrow did not claim that redundancy should be avoided, and HRO researchers did not claim that HROs never fail. Rasmussen noted that redundancy is difficult to manage, and that recent large-scale accidents were caused by systematically letting the system drift outside the design envelope. The willingness to pay for redundancy directed at very rare events might decline dramatically in periods of high competitive pressure. However, redundancy is essential for the operation of high hazard systems. It is not feasible to eliminate human errors. Over time, operators will explore the boundaries of safe operations, either deliberately or inadvertently.[15] Moreover, an organisation *needs* a certain frequency of reports on failures and incidents in order to validate the design assumptions and risk

[15] In Section 7.2 we will discuss how systems tend to drift toward the boundary of acceptable performance when faced with conflicting objectives.

predictions and to support risk management. It is thus difficult to see how a high hazard system can be managed without redundancy.

There is another contrast between Normal Accident Theory and HRO theory, which has received little attention. Proponents of HRO theory tend to assume that a single set of mechanisms can account for organisational reliability, across differences in technology, processes and organisational environments. In contrast, Normal Accident Theory postulates that the means to achieve reliable performance depend on the properties of the socio-technical systems: Systems with complex interactions call for decentralised control, whereas tightly coupled systems call for centralised control. We are not aware of empirical work directed at this issue.

5.9 Strengths and limitations of HRO theory

An important contribution of HRO theory is to direct attention to organisations with remarkable safety records and provide new insights into the functioning of these organisations. The dominating research approach is case studies of a single organisation or a few organisations. This approach does not allow researchers to isolate causal factors in the manner of a laboratory experiment.[16] However, we believe that a case study approach is the only feasible way to study complex patterns of organisational functioning in depth.

Most of the "classical" studies in the HRO tradition was directed at military organisations (aircraft carriers, nuclear submarines) or organisations that are strongly influenced by military culture (nuclear power plants[17], air traffic control). The discipline associated with some of these organisations would be unacceptable in Scandinavian work environments. It is necessary to ask whether the performance of HRO requires a culture characterised by military discipline. An alternative hypothesis is that organisations can build redundancy and received highly reliable performance in different ways. Preliminary results reported by Rosness et al. (2000) suggest that some offshore production platforms may have built a considerable degree of organisational redundancy. It is also an interesting issue whether the rapid turnover of personnel on aircraft carriers is an advantage or a disadvantage with regard to establishing a culture of high reliability.

It is not possible to build an HRO culture through attitude change campaigns or behavioural training alone. People in HROs report errors and take initiatives to improve flawed procedures because they expect these actions to lead to improvements. HRO theory thus does not provide an inexpensive alternative to continuous improvements of technology and working practice.

Rosness et al. (2000) suggested that concepts from HRO theory may help us understand how downsizing processes and low staffing levels may influence the safety performance of an organisation. Too low staff levels may remove the instrumental conditions for building organisational redundancy. Outsourcing might threaten the cultural preconditions for organisational redundancy, since personnel from different organisations might lack the mutual trust and openness necessary to consult, check and correct each other.

[16] Showing that an organisation with safety performance has characteristic X (e.g. extensive organisational redundancy) does not constitute a proof that characteristic X is the *cause* of the excellent safety performance. However, even case studies may be used for hypothesis testing (Yin, 1994). Sagan's (1993) is an example of a hypothesis testing case study.

[17] U.S: nuclear power plants were to large extent staffed by personnel with a navy background (former nuclear submarine crews).

6 Accidents as a breakdown in the flow of information: Turner's theory of Man-made disasters

Most major disasters are usually perceived as "fundamental surprises" by the media as well as by the organisations involved. However, several precursors or warnings are nearly always identified on hindsight by the media or accident investigators. This paradox is at the heart of Barry Turner's theory of man-made disasters (Turner, 1978; Turner and Pidgeon, 1997; Pidgeon and O'Leary, 2000).

The essence of Turner's information processing framework is that a disaster is almost always associated with recognition of a disruption or collapse of the existing cultural beliefs and norms about hazards. In seeking for a theory of disaster, Turner find himself concerned, not merely with systems of physical events - instead, he is concerned with a larger system which includes not only physical events, but also the perception of these events by individuals.

6.1 Notion of root causes and immediate causation

In developing his theory, Turner used reports from 84 accidents, which he systematically studied. He studied three serious accidents in depth. One of these was the Aberfan disaster in 1966, where a portion of a colliery tip on a mountainside slid down into a village and killed 144 people.

The Man-made disaster model proposes that accidents or disasters develop through a long chain of events, leading back to root causes like lack of information flow and misperception among individuals. Turner labels this chain, or time before a disaster, as "the incubation period". This is a developmental process where chains of discrepant events develop and accumulate unnoticed. This, Turner agues, is a result of a culture where information and interpretations of hazard signals fail. A typical accident can be traced back to initial beliefs and norms, which do not comply with existing regulations. Disaster development should be viewed as a process, often over years, developing from an interaction between the human and organisational arrangements of the socio-technical systems.

The incubation period starts with rigidities of belief and misperception of danger signals; events happen unnoticed or are misunderstood. Then, if someone takes action to the signals, it often results in what Turner label "the decoy phenomena". This is action taken to deal with a perceived problem which, on hindsight, is found to distract the attention form the problems that actually cause trouble. In many cases the company disregards of complaints from outsiders and fails to disseminate and analyse pertinent information. At the same time, the situation is not getting better when individuals often become insecure because of "out of date"-regulations and procedures, making the situation even more ambiguous.

Turner is not so concerned about immediate causation, but emphasise the breakdown in the flow and interpretation of information, which is linked to the energy or physical events. The critical assumptions in his theory concern the process leading up to disasters.

In the man made disaster model we also find stages after the actual disaster, including rescue and a final stage of full cultural readjustment to the surprise associated with the event. The whole model can be summarised in a simplified illustration (Figure 5):

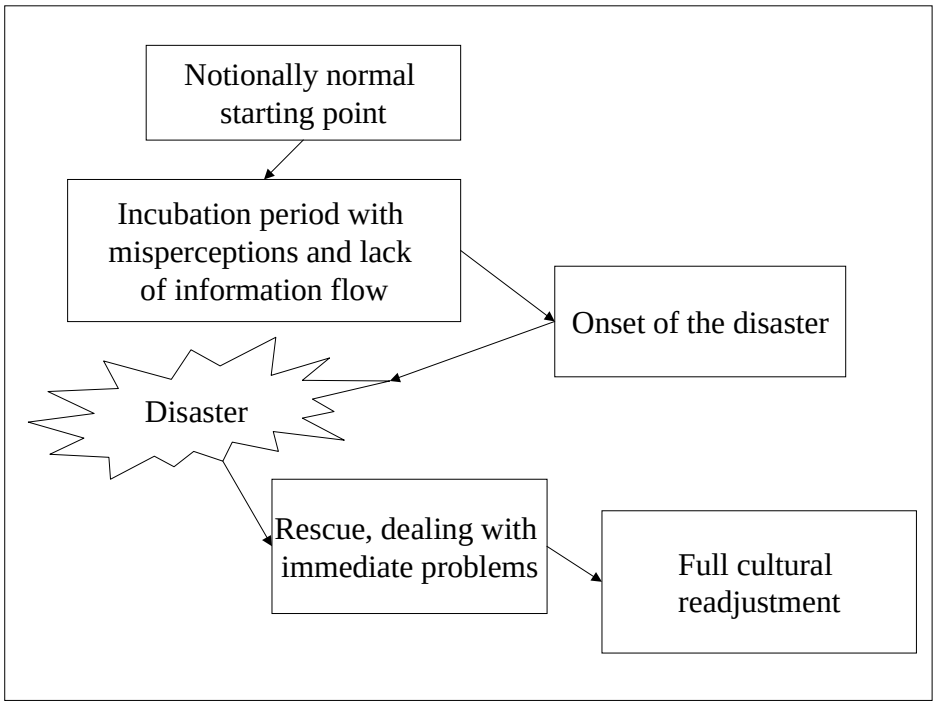

Figure 5. Main stages in the Man Made Disaster model of Turner (1978; Turner and Pidgeon (1997).

6.2 Cultures with requisite imagination

Ron Westrum's (1993) notion of cultures with requisite imagination nicely complements Turner's theory of Man-Made disasters. The expression "requisite imagination" is a paraphrase of Ashby's notion of "requisite variety". In its most compressed form, the law of requisite variety states that "only variety can destroy variety" (Ashby, 1981: 106). For an organisation to gain control over system it must be able to take as many distinct actions as the observed system can exhibit.

Westrum notes that organisations are very different in their ability to react to problems. He proposes the following key criterion of successful information flow in organisations (1993: 402):

> The organization is able to make use of information, observations or ideas wherever they exist within the system, without regard for the location or status of the person or group having such information, observations or ideas.

The variety in how organisations treat information is summarised in the typology shown in Figure 6. According to this typology, pathological organisations actively suppress warning, innovations and bridging, whereas these are actively promoted in generative organisations. Westrum cites several examples of projects where undesirable outcomes can be related to suppression of information.[18]

[18] An extended version of Westrum's typology is included in the TRIPOD safety management scheme, which was devised by researchers at The University of Leiden and the University of Manchester, and in the closely related "Heart and Minds" safety management scheme implemented by Shell International. The extended version includes five categories of organisations: pathological, reactive, calculative, proactive and generative.

Figure 6. A typology of how organisations treat information. Adapted from Westrum (1993).

Pathological	**Bureaucratic**	**Generative**
Don't want to know	May not find out	Actively seek information
Messengers are shot	Listened to if they arrive	Messengers are trained
Responsibility is shirked	Responsibility is compartmentalized	Responsibility is shared
Bridging is discouraged	Bridging is allowed but neglected	Bridging is rewarded
Failure is punished or covered up	Organization is just and merciful	Inquiry and redirection
New ideas are actively crushed	New ideas present problems	New ideas are welcome

6.3 Emergency plans as fantasy documents

Can an emergency plan contribute to misperception of danger signals? Clarke and Perrow (1996) claimed that organisational planning can produce "fantasy documents" which an organisation can come to believe in, to the extent that they ignore the bulk of experience showing that these fantasy documents may be inaccurate. They backed their claim by a study of an evacuation plan devised by the Long Island Lighting Company (LILCO). The plan should ensure rapid evacuation of a major area of Long Island in case of a catastrophe at the Shoreham Nuclear Power Station. Parts of the plan were put to test through a series of real-time exercises. The most extensive test was judged a success by LILCO itself as well as by the Federal Emergency Agency, whereas a three-judge panel threw the quality of the emergency organisation's preparations into doubt. Clarke and Perrow shared the latter view, and listed several rather serious problems that occurred during the exercise. However, they argued that the failure was not due to problems specific to LILCO, such as incompetent management, lack of preparation and expertise or lack of commitment. Rather, they launched a more general claim that *plans used to justify increasingly complex systems are often wildly unrealistic and they can impede organisational learning*. There is often no relevant historical record that may function as a reality check. Many accidents may not be covered by the plan. The plans are designed to be maximally persuasive to regulators, lawmakers and opponents of the system, and therefore tend to make benign assumptions about the environment. A bureaucratic emergency organisation with long lines of communication and excessive spans of control[19] may look impressive on the paper, but is not likely to produce an effective response in an emergency.

Clarke and Perrow claimed that such "fantasy documents" normalise the danger associated with complex, highly interactive systems by allowing organisations to claim that the problems are under control. This claim provides a possible link between the Normal Accident perspective and the information processing perspective on organisational accidents.

[19] "Span of control" refers to the number of subordinates that a superior is responsible for, as well as the variety of functions those subordinates must fulfil (Clarke and Perrow, 1996). Large spans of control (e.g. 10 – 20 subordinates) can be acceptable in routine jobs that are easily monitored, of similar function, and not interdependent.

6.4 Risk control strategies

Turner also proposed some strategies to control risk and prevent disasters. First, Turner stresses the significance of information flow when discussing risk control. It is important to be aware of some usual phenomena in organisations:

1) Completely unknown prior information: Where the information which foretells disaster is completely unknown, it is clear that there is little that can be done, except searching for better procedures for information flow in the relevant arena. We are still unlikely to experience such situations today; there is always someone who (should) know something relevant.
2) Prior information noted but not fully appreciated: Where information is potentially available, but not fully appreciated. The situation indicates that the information may not have been understood completely because individuals have a false sense of security when faced with danger signals. Often this emerges from distractions or pressure of work, which can give the subject an impression of the information as irrelevant.
3) Prior information not correctly assembled: When information about danger signals is carried in minds of individual humans, others can't reach it. A key to prevent disaster is therefore to place information in places where everybody can reach it.
4) Information available to be known, but which could not be appreciated because of conflict with prevailed understanding: In cases of disaster, Turner saw that relevant information was available, but when it was in conflict with prior information, rules or values, it was neglected and not taken into discussion.

To control risk, these "irrational" events have to be continuously evaluated by the organisation. A key factor is to make intensive efforts to collect and analyse information about hazards and find out what we do not know. Experiences from man-made disasters have shown that someone, somewhere do actually know something. The outcome of risk control therefore depends on the quality of monitoring risk.

Westrum (1993) discusses what can be done to develop organisations with requisite imagination. The organisation should provide incentives for thought. The only valid incentive for thinking is to use people’s ideas – and to make sure they know that their ideas are used. The organisation also needs to cultivate and reward efforts to bridge the boundaries between organisational layers, departments, subcultures and different sites. “Pop-out programmes” encourage the person with ideas and concerns to share them effectively. Pop-out programmes may, for instance, encompass the institution of new channels for information flow, empowering people to act when they see something that needs correcting, establishing open fora where workers can meet top-managers face to face and air complaints. It should be realised that pop-out programmes will only work if the organisation has the resources to act on the ideas and concerns that emerge from the process.

Many of the practices described in Section 5.6 are also relevant in an information processing perspective.

6.5 How can major accident risks be monitored?

The recurrent pattern of administrative and human failure, coupled with misinterpretation of warnings of disasters suggests, might, in theory at least, be identified through a holistic approach, to safety auditing (Turner and Pidgeon, 1997: 185). Detailed findings from the auditing process should be put together, in order to assess how the organisation handles information related to its vulnerabilities. We would need to identify indicators of the developing incubation period. Pidgeon

argues that existing technical hazard audits, such as Hazard and Operability Studies (HAZOPS), might be extended to incorporate relevant aspects of human and organisational failures.

There is the difficulty of unambiguously defining good and poor performance, which may be highly dependent upon the context within which an activity actually arises. Pidgeon notes that a number of the discussions of safety auditing share common ground with Total Quality Management (TQM). He therefore considers TQM a promising tool to control and audit risk.

Pidgeon also emphasises the role of safety culture as a key to handle and continuously monitoring risk (Turner and Pidgeon, 1997: 187-189). He argues that a good safety culture might both reflect and be promoted by at least four facets:
- Senior management are committed to safety
- Shared care and concern for hazards and their impact upon people
- Realistic and flexible norms and rules about hazards
- Continual reflection upon practice through monitoring, analysis and feedback systems

At the same time in Pidgeon emphasise the role of organisational learning to help initiating better risk perception among individuals, and by this overcome poor beliefs, norms and information flow in organisations (Turner and Pidgeon, 1997, 191-195). He emphasises the need for so-called double-loop learning (Argyris and Schön, 1978). It is not enough to change behaviour in response to feedback. We also need to improve out procedures for gathering and assessing signals about hazards, and to challenge our theories in use for interpreting the world.

6.6 Information processing related to the Åsta accident

One of the factors that made the Åsta accident possible was the absence of an Automatic Train Control system (ATC) on the Røros line. An ATC system would probably have stopped the northbound train within the station area of Rudstad, and the signalling system would have ordered the southbound train to stop before it entered the station area.

The initial plan for installation of Centralised Traffic Control on the Røros line included installation of ATC. However, the funding for ATC on the Røros line in 1993 was reallocated to other purposes. Repeated reallocations were made the following years, with the result that ATC installation did not start before 1999.

The absence of ATC on the Røros line violated the stated policy of the Norwegian State Railways, which was that all lines with remote control of signals should be equipped with ATC by 1995. The traffic safety manager warned about remote controlled sections without ATC at two top management meetings in NSB where the managing director was present in 1995. In 1996 the traffic safety manager issued a memo where he repeated his concerns. In 1997 he repeated his concerns in a new memo.

The Commission of inquiry asked the managing director of NSB at that time about his knowledge about the safety manager's concerns (NOU 2000:30, p. 153, our translation):

> The managing director of NSB at that time, Ueland, could not remember that he had received [the memo from the traffic safety manager. He explained to the commission that he was confident that the consequences of [not giving priority to ATC installation] had been assessed, and said that there had been no disagreement in the organisation about the reordering of priorities. He claimed that nobody in the organisations had said that the priorities could not be changed, and that one could not postpone installation of ATC any longer. He further claimed that it was a clear judgement in the organisation that they had a safe and good system. On a question from the Commission about whether he

> considered the safety on the Røros line on the 4th of January 2000 adequate, Ueland explained that the issue of safety was simple to him; it was either safe to drive trains, and then the trains would roll, or otherwise the trains stood still. He claimed that he, like many others, had been living in the belief that it was safe to drive on the Røros line.

This excerpt illustrates the paradox which inspires the information processing perspective. The knowledge about the problem exists somewhere in the organisation or its close environment, but this knowledge is not shared by the dominant decision makers, and therefore not acted on.

How did this situation come into being? The safety director at the time accepted the first reallocation of funding, since it implied a delay for one year only. However, his follower was placed at a position lower down the organisational hierarchy, and thus had more difficult access to the attention of the top management. At the same time, the confidence was developing in top management that the organisation tackled its safety challenges well enough to concentrate on other issues, such as punctuality and the development of new services. This conviction was founded on hard data – a favourable long time trend in fatal railway accidents, culminating in two very good years (1996 and 1997). There existed no strong external "watchdog" who could effectively challenge this conviction[20]. This pattern fits well into Turner's notion of an incubation period where the organisation systematically disregards warning signals. It also suggests that power relations play an important role in organisational information processing.

6.7 Strengths and limitations of the information perspective

An important contribution of the information perspective is Turner's finding that during the incubation period, there is nearly always someone who is aware of the imminent danger. This finding has strong implications for safety management: The accumulation of more data *per se* does not prevent accidents. It is necessary to focus on the processes through which information is disseminated, combined and interpreted. We need to understand the mechanisms through which some warnings gain the attention of decision-makers and eventually lead to preventive action, whereas other warnings are ignored or rejected. The research challenges raised by this finding are far from resolved. It may prove necessary to go beyond a narrow information perspective and, for instance, explore whether a political or power perspective is also needed to get a grip on these phenomena.

A fundamental challenge for researchers applying the information perspective is to *show that their claims are meaningful and valid to actors who do not have the benefit of hindsight*. This can be illustrated by an example. Turner reported that organisations often fail to take action on danger signals because "decoy phenomena" distract the attention from the "real" danger signal. Was it really possible to distinguish between "decoys" and warnings that would materialise *before* the accident happened? Or do we need the information provided by the accident to be able to label some warnings as "decoys" and others as "real"? If this distinction can only be made based on hindsight, then this finding is of little help to persons charged with preventing accidents. When reporting findings based on hindsight, researchers should therefore strive to take the perspective of actors who do not have the benefit of hindsight, and ask themselves whether the finding still makes sense.

[20] NSB was thus not subject to external regulation of safety before the Norwegian Railway Inspectorate was established in 1996. The Inspectorate had very limited resources in its first years (NOU 2000:30).

7 Risk handling in the face of conflicting objectives: Risk taking, adaptation and drift

We live in an open market economy. This economy is "designed" to exterminate organisations that use more resources than absolutely necessary to deliver a given product. Organisational survival is thus a matter of balancing on the edge. Risk control and safe performance often requires considerable resources such as money, time, and competent personnel. Humans or groups may make risky choices when facing a dilemma. Moreover, performance at the level of individuals, groups and larger organisational units may drift over time under the pressure of conflicting objectives. It is thus impossible to give a balanced view on organisational resilience without considering how organisations handle conflicting objectives.

7.1 Taking a risk or running a risk?

What exactly happens when people face the choice between a risky course of action and a less risky course of action, and act in a manner that eventually triggers an accident or fails to recover a dangerous situation? Very often, they do not *face* the choice at all. In a study of 57 accidents at sea, Wagenaar and Groeneweg (1987) investigated whether the negative outcomes were the result of deliberate risk-taking, i.e. whether the captain deliberately selected a dangerous course of action, knowing that there existed feasible and less risky alternatives. They found that this was the case in only one of the 57 accidents.[21] In 21% of the cases the information of the immanent danger was not even available, and in another 27% the situation was not recognised as problematic. In a further 36 % the consequences were either not foreseen, or the likelihood of disaster was underestimated. Wagenaar and Groeneweg thus claimed that a large majority of the captains had been completely taken by surprise. They had been *running a risk* rather than *taking a risk*.

This result should not be uncritically generalised to all settings were risk-related decisions (or "non-decisions") are made. Personnel at the "sharp end", those who work close to the sources of danger (e.g. process operators, train drivers, pilots, captains), face very strong incentives to avoid accidents. They and their fellow workers may risk their lives if an accident should occur. They are also particularly susceptible to blame, since the causal chain between actions at the sharp end and the unwanted consequences is usually short and conspicuous.

Things may be different at the "blunt end" – for instance at the administrative quarters of shipping companies. After the capsize of the Harold of Free Enterprise, the company management turned down several applications for an alarm on the bridge to prevent the ship from leaving harbour with the bough doors open, despite the fact that this had happened to a sister ship. Several other requests to correct urgent safety problems had been handled in similar manner.[22] In this case, company management was informed about the risks resulting from their decisions, and alternative courses of action had been proposed. It is thus a reasonable claim that these managers were deliberately *taking a risk*. High level managers may be more prone to taking risks for two reasons (Rasmussen, 1994a). Due to their professional background (e.g. business schools) and their distance from daily operations, some of them may fail to fully comprehend the implications of the

[21] In that case the captain claimed that the risk of a collision was smaller than the risk of grounding, a judgement which the Dutch Shipping Council decided had been wrong.

[22] Department of Transport/1987): mv Herald Of Free Enterprise. Formal Investigation. Report of Court No. 8074. London: Her Majesty's Stationery Office.

warnings they receive from the sharp end. Secondly, the incentive systems of many managers may direct their attention to short term profits at the expense of the prevention of adverse events that they perceive as unlikely to happen during the few years before they move to another position.

7.2 Migration of activities towards the boundary of acceptable performance

Rasmussen (1997) suggested that we may think of the handling of conflicting objectives in terms of activities migrating toward the boundary of acceptable performance (Figure 7). The basic idea is that human activities are characterised by continuous adaptive search in the face of partially conflicting pressures and needs. Individuals and groups strive to keep the workload at a comfortable level, to find some intellectual joy in the activity, and to avoid failure. They face requirements and pressures with regard to e.g. productivity and quality. In practice, the work place allows them considerable freedom to try out different ways to handle these partially conflicting needs and constraints. This is depicted as the space between (1) the boundary of financially acceptable behaviour, (2) the boundary to unacceptable workload, and (3) the boundary of functionally acceptable behaviour with regard to risk. In seeking for a viable adaptation, humans will explore this “work space”. Rasmussen compared this to the brownian movements of molecules in a gas. Both the “effort gradient” and the “cost gradient” are likely to drive the activities towards the boundary of safe (i.e. functionally acceptable) performance. An error or accident may occur if the crossing of this boundary is irreversible.

This model directs our attention to what happens at the boundary of safe performance. Management may seek to define a narrower boundary through safety campaigns, and thus increase the safety margins in the activities. We may also ask what happens when an activity approaches or crosses the boundary? Will the actors receive an insistent warning from the system and have the opportunity to reverse their actions? Since many dangerous situations do *not* lead to disaster, is there a risk that they will adapt to warnings over time? May they even modify their mental models of the system in such a way that they ignore the dangers involved in crossing the boundary of safe performance?

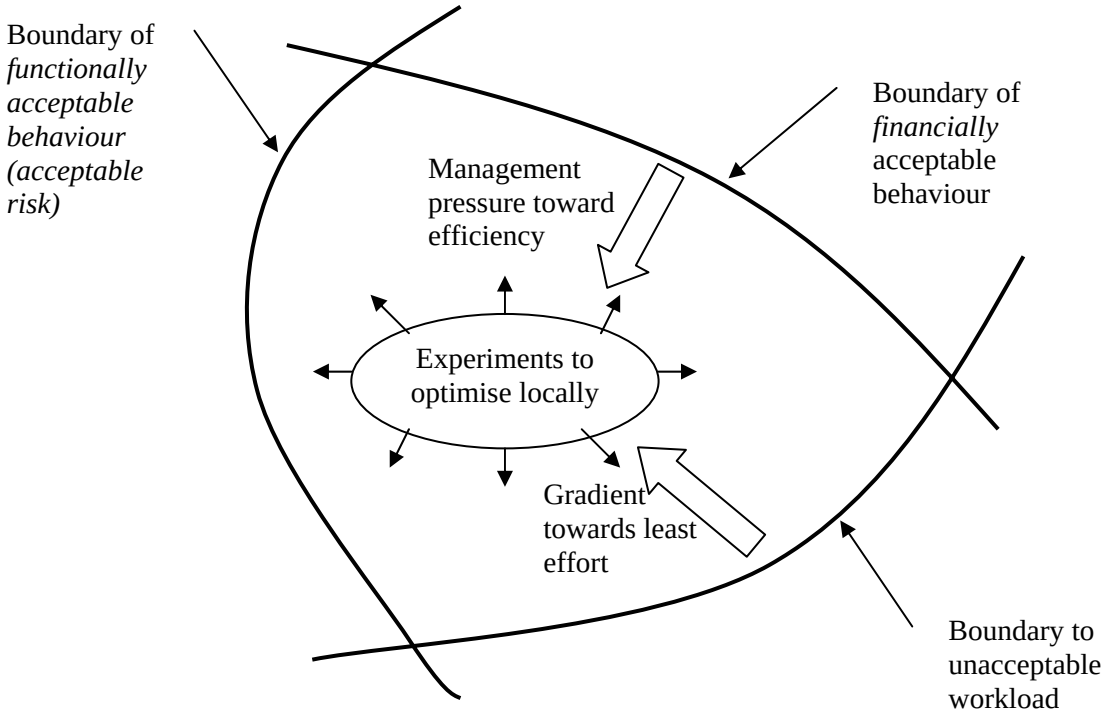

Figure 7. Under the pressure of conflicting objectives activities tend to migrate toward the boundary of acceptable performance (Adapted from Rasmussen, 1996).

7.3 Distributed decision making

In a complex system, many activities take place in parallel. At a given moment, each actor may have incomplete or inaccurate knowledge about the state of the system and the ongoing activities. Moreover, the parallel activities may interact in non-obvious manners if the system is characterised by high interactive complexity (see Section 4.3). A system is characterised by *distributed decision making* to the extent that it lacks a centralised decision-maker and each decision-maker has a model and information of a limited part of the problem (Brehmer, 1991)[23].

The migration model as presented in Figure 7 above captures a single activity performed in isolation. Figure 8 illustrates the adaptation process in a complex system with distributed decision making (Rasmussen, 1994b). Actions within one activity may change the boundary of acceptable performance for another activity. For instance, on an offshore production platform, one work team may open a valve on the flare system to drain off liquids before they replace a valve. Another

[23] Distributed decision making differs from group decision making, where the problem is to achieve consensus when everyone is capable of understanding the whole problem.

work team at a different place may at the same time need to release hydrocarbons into the flare system in order to start a their job. However, due to the first activity, the second work team cannot safely perform a normal pressure release. The figure therefore contains two different boundaries for acceptable (i.e. safe) performance. The inner boundary delimits the *unconditionally* acceptable state of affairs, i.e. acceptable without regard to the behaviour of the other actors. The outer boundary delimits the *conditionally* acceptable state of affairs. A single actor can enter the area between the two boundaries without unacceptable risk of triggering an accident. However, the risk may become unacceptable if two or more actors enter the area between the two boundaries. Rasmussen (1994b: 28) termed such coincidences 'singularities'. They are characterised by dramatic shifts in system performance and are often perceived as "basic surprises".

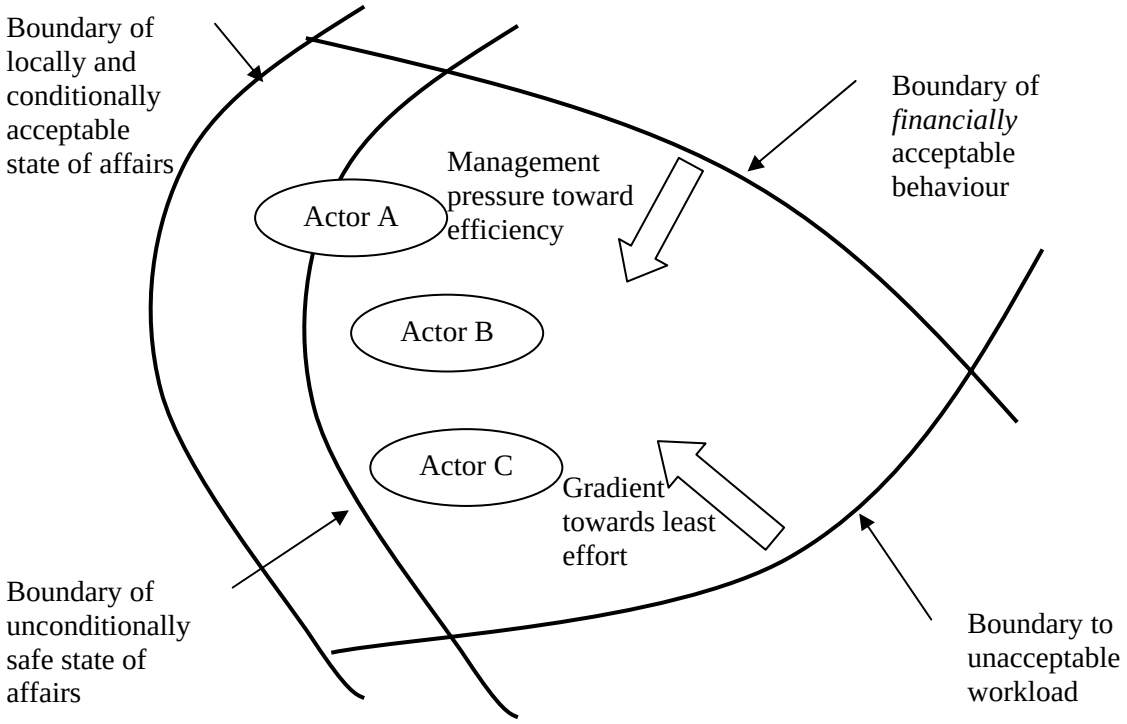

Figure 8. Adaptation in a complex organisation, where several actors are migrating more or less independently within the space of acceptable performance. (Adapted from Rasmussen, 1994b).

Again, the crucial issue is how the actors behave with regard to the two boundaries. In this case, we may expect actors to strive for *local optimisation*, based on their incomplete knowledge about the system. They will take into account the dangers and potential scenarios they know about, but not those that are not "visible" from their local point of view. The implication is that actors will run risks and singularities will occur, unless the inner boundary is made active and able to bound the natural migration towards the outer boundary. Rasmussen (1994b) suggests that it is feasible to provide the necessary decision support to help operators stay within the unconditionally safe boundary in a well-structured process plant. Providing visible margins to safety boundaries may even increase operations efficiency, since the operators will not need to maintain an excessive margin to an invisible boundary.

Distributed decision making is an answer to the complexity of the problems and fast pace of changes in the decision contexts facing the decision-makers. Moreover, the degree of centralised control may be reduced as a consequence of efforts to reduce management and administration costs. This sometimes happens through explicit organisational change. At other times, the reduction of centralised control may be unplanned and inconspicuous. Managers may be given new responsibilities and time-consuming tasks, and adapt by reducing the attention they give to supervisory tasks (e.g., checking that work permit forms have been adequately completed).

Many organisations have developed administrative systems in order to manage the risks associated with parallel activities and distributed decision making. For instance, work permit systems are used in the process industry to make sure that critical tasks are properly co-ordinated, and that necessary precautions are taken. These administrative systems may be even more safety-critical than many technological barriers, because some tasks involve the temporary removal of several technical barriers. A failure related to the work permit system might thus hit the system in a very vulnerable state.

7.4 Levels of decision-making

Yet another dimension of decision- making remains to be explicated. The control of risk, as well as the production of accidents, takes place at many levels, ranging from political systems to individual operators and even technical systems (e.g., automatic process control and safety systems). This is illustrated in Figure 9. Each level can influence each other in an integrated and tightly coupled system. Higher levels can influence lower levels through, e.g., explicit instructions, by the provision and limitation of resources, by establishing incentive systems, or by determining *how* decisions are to be made at lower levels. On the other hand, lower levels may use discretion when they interpret and implement directives from higher levels, they may control the information flow to higher levels, or they may bypass a level and direct a lobbying effort at the level above.

The figure also illustrates the traditional formal means used by each level to control the level(s) below. Rasmussen (1997) argued that the classical prescriptive command-and-control approach, where rules and conduct are derived top-down, works far too slowly to tackle the dynamics of modern economies.

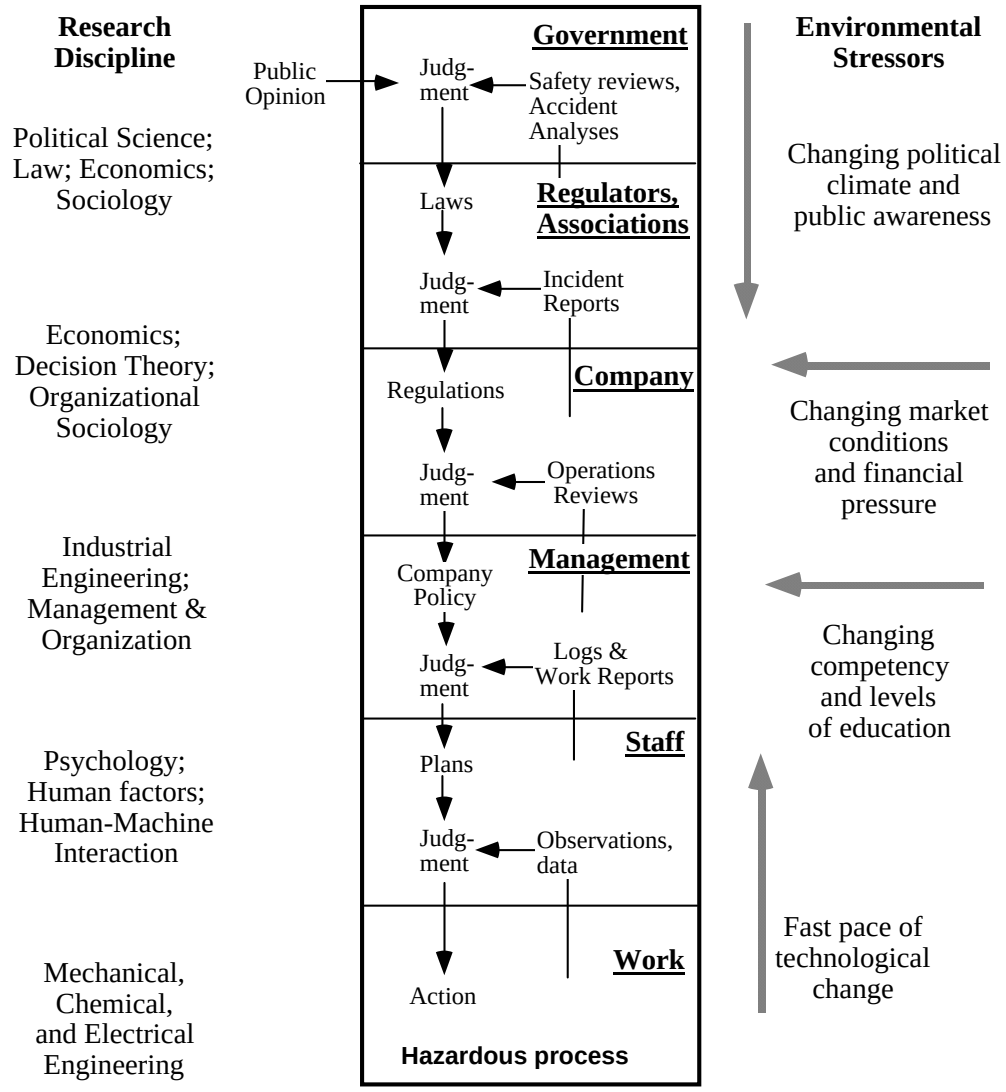

Figure 9. The socio- technical system involved in risk management (Adapted from Rasmussen, 1997).

7.5 The diversity of decision contexts and decision processes: A contingency model

We have hinted at the diversity of decision settings in earlier sections in this chapter. We will now present a typology, which may make this diversity more comprehensible.[24] We have noted that some decisions are made at the "sharp end", i.e. close to the hazard sources. Others are made at the blunt end, removed from the hazard sources. Decisions makers and decision settings also differ in their level of authority. A manager can issue orders and directives to his or her subordinates, and an inspectorate can issue directives to and impose sanctions on companies. We can thus characterise decision settings according to two dimensions, (1) proximity to the hazard source, and (2) level of authority. This is illustrated in Figure 10. Pilots, offshore platform superintendents or aircraft line maintenance personnel usually find themselves at the sharp end, i.e. close to the hazard source. Designers, planners, analysts and regulatory institutions typically operate at the blunt end. Some actors may be "operationally" close to the hazard source, even though they are physically remote, for instance air traffic control operators or centralised train

[24] The typology was introduced by Rosness (2001) and has been elaborated in an unpublished draft paper by Rosness and Hovden and in a paper by Kørte et al. (2002).

control operators. We will consider these actors as belonging to the sharp end, even though they are less vulnerable in case of an accident. Actors at the sharp end are often mostly event-driven and thus operate within a shorter time horizon most of the time. We also expect actors at the sharp end to have more updated and detailed hands-on knowledge of the system they operate than actors at the blunt end.

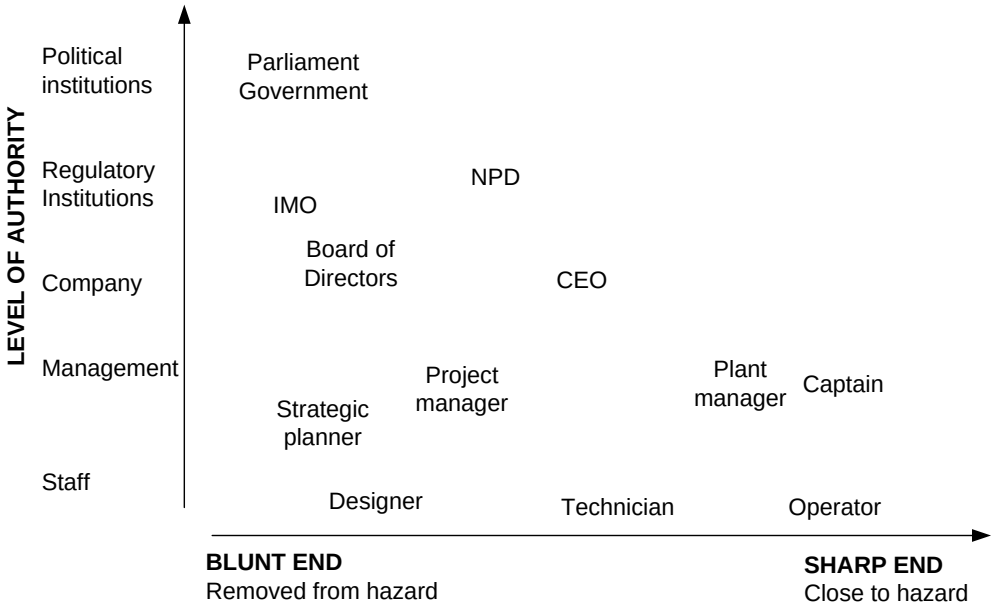

Figure 10. Two dimensions for characterising setting for safety related decision making, adapted from Rosness (2001). (IMO - The International Maritime Organisation; NPD - The Norwegian petroleum Directorate; CEO - Chief Executive Officer.)

The conditions under which actors make decisions strongly influence the decision processes and outcomes. We thus expect decision criteria, procedures, and outcomes to be related to (1) how close an actor or decision forum is to the hazard and (2) the level of authority of the actor or forum. These relationships are complex, since decision-makers also adapt to circumstances not covered by these two dimensions. However, we believe that even a grossly simplified model of these relationships may be helpful in sensitising us to the way decision-makers adapt to their setting. We therefore identify five distinct decision settings and propose an associated typology of decision modes, see Figure 11 and Table 4.

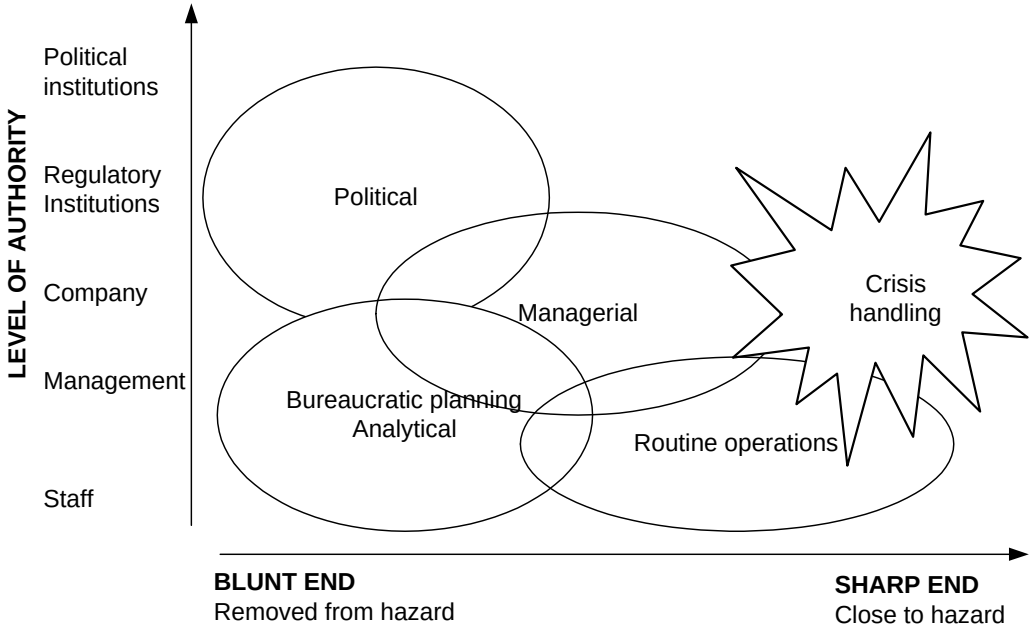

Figure 11. Classes of decision processes. Adapted from Rosness (2001).

In order to illustrate the logic of the model, we will consider *routine operations* in some detail. Actors at a low level of authority and close to the hazard, such as drivers, process operators, and ship crews, often experience uneven workloads because their tasks are event driven. Their decision making is often constrained by limited situation awareness (Woods et al., 1994). They may not receive the requisite information to build a complete and updated model of the situation, they may not have enough free information processing capacity to maintain an updated system model, or they may lack mental models that adequately represent the properties of the larger system. Actors in complex systems are likely to resolve goal conflicts within a condition of local rationality. He/she is not in a position to assess the overall impact of their choices, or to assess how their choices interact with those of the other actors (Brehmer, 1991; Rasmussen, 1997).

We propose that sharp end actors at low levels of authority tend to focus on smooth and efficient operations. At the same time they try to keep the workload at a comfortable or at least tolerable level. Human behaviour in routine tasks is to a great extent automated and feed-forward driven, with occasional, progress checks, references to rules or reactions to signals from the environment (Rasmussen, 1986; Reason, 1990). This mode of behavioural control allows very efficient and smooth performance. However, due to the limited attention and feedback control, highly automated behaviour is prone to slips ("actions not as planned"). The reaction to danger signals may be automated to a point where the signal escapes conscious processing. Operators' mental models of the system may decay over time because some aspects of the models are not maintained through their daily working experience. Safety margins may erode over time if the system fails to provide clear warnings if the operator performance approaches or crosses the boundary to unacceptable risk or if the operator adapt to these warnings.

The example also illustrates the gross simplifications inherent in this model. The work of sharp-end operators at low authority levels may include significant amounts of non-routine work, for instance concerned with the handling of irregularities in the production system. The share of non-routine tasks may be increasing in many production systems, as routine tasks become automated. One might think of a continuum of decision modes, ranging from "pure" routine operations" via

skill-based, rule-based and knowledge-based problem handling to the handling of crisis situations with imminent danger of a catastrophic outcome.

Table 4. Dominant constraints, decision criteria and typical problems in different decision modes.

Decision mode	Dominant constraints	Dominant decision criteria	Typical problems
Political	Conflicts of interest	Robust consensus	Inconsistency Non-optimal decisions Erosion of safety margins
Managerial	Information processing capacity	Find an option that is good enough (satisficing)	Inadequate problem definitions Stick to SOP Erosion of safety margins
Analytical & Bureaucratic	Hands-on knowledge	Comply with rules & standards Optimise selected attributes	Unrealistic assumptions Deficient models Erosions of safety margins
Routine operations	Workload Situation awareness	Smooth, efficient operation Optimise workload	Slips Miss warnings Local rationality Erosion of safety margins
Crisis handling	Stress Time to obtain information and act	Avert catastrophic outcomes Avoid extreme stress levels	Defective coping if danger materialise

As summarised in Table 4, we may expect actors in other types of decision settings to face other constraints, and to adapt by focussing on different decision criteria. On this basis, we may make informed speculations on the types of problems that may be related with the outcomes of decisions. More comprehensive presentations of the model are given in Rosness (2001) and in Kørte et al. (2002).

7.6 Adherence to rules, culture and resources

We have all heard that rules are there to be broken. It has also become a commonplace in discussion on organisations to point to the gap between espoused theory and theory in use (Argyris and Schön, 1978), between what people say they do and what they actually do. Is it then a law of nature that any rule or procedure will be bent or broken as soon as it meets the harsh reality of conflicting objectives and rapidly changing environments?

This issue brings us back to research on High Reliability Organisations. Bourrier (1998) studied the daily adjustment and modifications made to maintenance procedures and work orders in two French and two U.S. nuclear power plants. The need for adjustment may arise from unplanned situations and changes, wear and tear problems, or the availability of tools.

In one French power plant, she found that workers did make minor adjustments, which they did not report to management. The adjustments were not totally out of control, since the maintenance staff shared among themselves a set of tacit rules that were different from the formal rules. These rules were conveyed to new workers through a socialisation process. However, upper management was left with very limited knowledge about the exact way in which the problems

were handled. In the two U.S. power plants, however, Bourrier found that workers were strictly following the rules, asking their foremen for help and guidance or new work orders each time they ran into an unplanned difficulty.

It is tempting to explain this contrast by referring to the national cultural clichés, e.g. rigorous American contractual orientation versus a pragmatic French habit of getting along and muddling through. However, Bourrier found that an alternative explanation could be based on organisational differences. She examined the opportunities for bridging the gap between procedures and practice. At the French plant, no resources were provided for adaptation of formal procedures to unforeseen situations. The personnel authorised to revise procedures were not available to the maintenance staff. Compliance with procedures was thus not an option if the job was to be done. The maintenance personnel compensated by developing and propagating their own norms, for instance tolerances that were somewhat larger then those accepted in the formal procedures. Such unofficial norms were saved in personal notebooks and conveyed to trainees, but they were not known to plant management or to the engineers who wrote and revised maintenance procedures.

One of the US plants (Diablo Canyon) bridged the gap between procedures and practice by making the engineers responsible for procedure updates available to maintenance personnel. The workers responded by taking all problems to the engineers. It was part of the worker culture at this plant to take very few initiatives, since managers and engineers were paid to do the thinking. At the other US plant (North Anna) maintenance foremen had the formal authority to initiate procedure adjustments and updates. They also had the time needed to handle these problems.

The issue of compliance thus has two sides. It is not only a matter of controlling operator actions. This was to some extent achieved in all three power plants. The crucial difference is that the US power plants functioned as self-correcting organisations. They were able to develop explicit (as opposed to tacit) mechanisms to reconcile formal norms with the realities of daily work. However, this learning does come at a cost. At Diablo Canyon, the operating costs were very high due to the staffing level. The US plants also had better availability factors than then the French plant.

7.7 Implications for risk control and risk reduction

The organisations studied by LaPorte and Consolini (1991) had abundant resources to fulfil their strategic missions and to prevent catastrophic failure. They were not under the pressure of a competitive market. Some of the nuclear power plants (e.g. Diablo Canyon) had a very large and complex organisation (Schulman, 1993), and could allocate a lot of time and engineering capacity to the planning of maintenance shutdowns (Bourrier, 1998). However, such abundance is the exception rather than the rule. Building an enormous organisation is simply not an option in most industries. Conflicting objectives are here to stay, and the safety discipline has to find ways to cope with scarcity and dilemmas.

A common strategy for handling conflicting objectives is to create independent institutions or organisational units (“watchdogs”) to regulate or monitor safety performance (e.g. Lindblom, 1959). For instance, Lord Cullen, in his investigation of the Piper Alpha disaster, pointed out the importance of the regulatory body being perceived as independent of commercial interests.

Some strategies for proactive risk reduction efforts can be derived from Rasmussen’s “migration” model (e.g., Rasmussen and Svedung, 2000; see Section 7.2 above):

- The boundaries of safe performance should be made visible and touchable. The actors should have a way to know when they approach or exceed the boundaries. Design envelopes should

be communicated effectively to operators. A friendly system will also allow them to recover in case he or she momentarily exceeds the boundary. This strategy is complicated by the fact that humans tend to adapt to warnings. We also need to bear in mind that human information processing capacity is limited. In many systems, the filtration of unnecessary alarms may be an effective measure to make the boundaries visible.

- Pressures that drive decision-makers toward the boundaries of safe performance can to some extent be counteracted. Managers may, for instance, by follow up safety performance on a par with economic performance.
- It is a good idea to communicate explicitly about trade-offs. Operators are put in a difficult double-bind situation if managers state that safety has priority, but tacitly communicates the opposite message through planning, follow-up, resource-allocation or their own example (Woods et al., 1994).

External pressures on organisations do not only cause dilemmas. Many dilemmas arise from high level decisions, for instance when insufficient time is allocated to the completion of a project. It is important that high level decision makers who are in a position to put lower level actors under pressure also are held accountable for the accident risks associated with their decisions.

The migration model suggests that regulations and procedures have two important functions with regard to safety. The first function is to help actors stay within the boundary of safe performance. The second function is to prevent conflicts between activities when decision-making is distributed. The first function does not always call for a detailed specification of the one and only acceptable way to perform a given task. The point is rather to help the actor identify the boundaries of safe performance. The second function calls for procedures, which highlight the interfaces with other activities, where there is a potential for conflict. In the latter case, a rather rigorous standardisation may be necessary to make the performance of other actors predictable.

7.8 Conflicting objectives and the Åsta accident

Based on the public investigation report, we may identify examples of conflicting objectives related to safety at several levels:

- The investment in Automatic Train Control had been postponed for years on the Røros line, in spite of several warnings. In the early nineties, funding was made available, but the organisation did not allocate sufficient planning resources and top management attention to ATC installation. In the later nineties, the project was given low priority in the budgets. ATC seems to have competed with preparations for the 1994 winter games, as well as the Gardermoen railway and the introduction of high-speed trains.

- In 1997, the Norwegian Railway Administration introduced a new departure procedure. The former departure procedure required the train guard and the driver to check exit signals independently. The new departure procedure was adapted to main lines equipped with ATC. Since a failure of the driver to obey a red exit signal would be recovered by the ATC, the conductor was not required to observe the exit signal. He could thus concentrate on the safety on passengers leaving and boarding the train. The problem was, of course, that no double-checking would take place on railway sections without ATC. It is not clear whether this was merely a conflict between two competing safety considerations, i.e. safety of embarking and disembarking passengers versus collision risk. The Åsta commission mentions that according to the new departure procedures, the Norwegian Railway Administration might permit passenger trains to operate without a conductor, thus reducing personnel costs.

We discussed the context of the decisions to postpone ATC installation in Section 6.6. We are now in a position to add another factor. Safety management in Norwegian railway operations has traditionally taken a reactive approach, with a focus on detailed operational rules (NOU 2000:142). Managerial action to improve safety was typically taken in response to incidents and accidents, whereas it was uncommon to systematically assess whether a proposed change might jeopardise safety, or whether a given activity or system was safe. This approach may have lead to a *lack of clear and compelling criteria to identify the boundaries for acceptable risk* in decision situations that are not covered by the operational rules. This problem is also illustrated by the way the managing director at the time commented on the safety of the Røros line. Note that there is no reference to criteria as to what is safe enough (NOU 2000:30, p. 153, our translation):

> He further claimed that it was a clear judgement in the organisation that they had a safe and good system. On a question from the Commission about whether he considered the safety on the Røros line on the 4th of January 2000 adequate, Ueland explained that the issue of safety was simple to him; it was either safe to drive trains, and then the trains would roll, or otherwise the trains stood still. He claimed that he, like many others, had been living in the belief that it was safe to drive on the Røros line.

These examples also illustrate *decoupling of decisions that are related in their impact on safety*. For instance, the decision to implement a new departure procedure was not coupled to the installation of ATC, although the new departure procedure might lead to increased collision risk on railway sections without ATC.

We have not identified a unified theory that gives a comprehensive grasp on the issues related to the handling of conflicting objectives. This chapter thus combines concepts and models from several authors. It may be too early to discuss the “strengths and limitations” of this perspective. In the future, we should not be satisfied just to note *that* some accidents may be related to the presence of conflicting objectives. Rather, we need to understand why some organisations handle conflicting objectives in a safer manner than others. This may enable us to propose *how* an organisation may survive in a competitive environment without unnecessarily compromising safety.

Safety scientists have been rather reluctant to use the concept of “power” until now. It may prove necessary to examine how power is built and exploited in safety-related decision in order to understand how conflicts between conflicting objectives are handled in a setting where different actors have different interests.

8 Summary and comparison of the perspectives

In the previous chapters, we discussed one perspective at a time. It may not be quite clear how much the perspectives have in common, or to what extent they are contradictory or complementary. We will therefore compare the perspectives with regard to a few central issues. The comparison is summarised in Table 5.[25]

8.1 Notions of immediate causes of accidents

The energy and barrier perspective provides an explicit view on the immediate causes of accidents: Accidents occur when objects are affected by harmful energy out of control, in the absence of effective barriers between the energy source and the object. This view of accident causation seems to be explicitly or implicitly accepted by most people that are occupied by accident prevention and risk analysis.

Proponents of the other perspectives do not explicitly contradict the energy view of accident causation. However, it is elaborated in different ways. In Normal Accident Theory, Perrow (1984) restricts attention to systems accidents, which are caused by unanticipated interactions of multiple errors. Researchers in the HRO tradition tend to say little explicitly about the nature of accident causation, but their implicit idea seems to be that accidents are triggered by errors that have not been recovered in time. Turner's theory of man-made disasters explicitly associates accidents with a breakdown in the flow and interpretation of information within an organisation. In the perspective of conflicting objectives, accidents may be seen as the results of actors transcending the operational envelope of the systems they operate. At this level, the perspectives are complementary, rather than contradictory.

8.2 Notions of "root causes" of accidents

With the term "root causes" we refer to system attributes or processes that are used to explain why the immediate causes of accidents occur. Proponents of the energy perspective tend to focus on failures to establish and maintain adequate barrier functions. They may also point to dependencies among barriers, since such dependencies may increase the likelihood that several barriers may fail simultaneously.

Normal Accident theory finds an underlying problem in a mismatch between the properties of the system and the control strategy. Such a mismatch occurs if a centralised control regime is applied to system with high interactive complexity, or if a decentralised control regime is applied to a tightly coupled system. A tightly coupled system with high interactive complexity is inherently

[25] We have not included a comprehensive discussion on research needs in this memo. However, a few areas may be mentioned: The research on resilient organisation has been conducted on aircraft carriers and in the nuclear industry. These types of industries have a lot of resources and military organisations do not have to consider profit to operate. We need more studies from other industries to examine whether the findings can been found in other areas. And it is particularly important to study industries with fewer resources than aircraft carriers and nuclear industry.
We need more field studies of decision-making in normal settings. What is actually happening in complex environments concerning decision making? Are organisations in the oil and gas industry able to reconfigure in response to emergencies or periods of extreme demand?
One important lesson after the terror attacks on September 11, 2001 is that the risk and safety community will have to pay more attention to the danger for terror attacks and sabotage. We need to assess the extent to which the perspectives outlined in this memo can be fruitfully applied to issues related to intentional harm.

vulnerable, since a mismatch is bound to occur, irrespective of whether a centralised or a decentralised control strategy is applied.

The theory of High Reliability Organisations does not say much explicitly about root causes of accidents. By turning around their findings from organisations with very reliable performance, one may suggest that less reliable organisations may be characterised by a failure to build organisational redundancy and by a failure to adapt the organisational structure in the face of demanding situations.

Table 5. Comparison of the perspectives on resilient organisations.

Issue	Energy and barrier perspective	Normal accident perspective	HRO perspective	Information processing perspective	Conflicting objectives, adaptation and drift
Notion of immediate causation of major accidents	Object affected by harmful energy which is out of control in the absence of effective barriers between energy source and object	System accidents characterised by unexpected interaction of multiple errors, some of which are usually latent	There seems to be an implicit understanding that accidents are caused by un-recovered errors.	A breakdown in the flow and interpretation of information which is linked to the physical events.	Actors cross boundaries towards unacceptable risk in effort to locally optimise behaviour.
Notion of root causes of major accidents	Failure to establish and maintain adequate barrier functions Dependencies among barrier functions	Mismatch between system properties (complexity, coupling) and control strategy. Contradiction between demand for decentralised control and centralised control in complex, tightly coupled systems	Not discussed explicitly.	Disasters develop as a process in which the prominent perceptions and interpretations in the organisation gradually diverge from the realities manifested by danger signals and warnings from the outside.	High level decision makers taking risks and running risks. Unforgiving systems (invisible and "un-touchable" boundaries). Distributed decision making in dynamic and opaque systems.
Critical assumptions	All accidents involve energy flow out of control as the immediate cause of harm. Effective risk control strategies can be found by focussing on energy flows.	Systems with high interactive complexity require decentralised control. Tightly coupled systems require centralised control. Organisations can not be centralised and decentralised at the same time.	"Human errors are here to stay". However, it is possible to achieve nearly faultless performance through organisational redundancy. Organisations can change and adapt to different situations.	A disaster is almost always associated with recognition of a disruption or collapse of the existing cultural beliefs and norms about hazards.	Activities tend to migrate towards the boundary of acceptable performance as actors' search viable trade-offs between such considerations as workload and productivity.
What is the relationship between minor and major accidents?	Minor and major accidents have the same basic causes. However, major accidents tend to involve failure of more than one barrier.	Minor accidents are often caused by a single failure. Major accidents are caused by multiple failures and are related to the structural properties of the system (complexity, coupling, and control).	Not explicitly discussed. In a HRO, one would expect major accidents to involve failure of one or more recovery mechanisms.	Many smaller accidents can be indicators for disasters (large-scale accidents), but they do not necessarily have the same root causes - failures in information processing.	Major accidents tend to arise through a pattern of distributed decision-making and conflicting objectives, more often than minor accidents do.

In Turner's theory of Man-made disasters, accidents are viewed as the culmination of a process in which the prominent perceptions and interpretations in the organisation gradually diverge from the realities manifested by danger signals and warnings from the outside. This process may often includes decoy effects, where a less important problem distracts the attention form the problems that actually cause trouble

In the perspective of conflicting objectives, several mechanisms may contribute to accidents. High-level decision-makers may deliberately *take a risk*, or they may *run a risk*, i.e. make a decision (or non-decision) which affects the risk level without considering the impact of that decision. The impact of such decisions may be indirect. For instance, a high-level decision-maker may fail to provide the resources that are needed for safe operation of a system. Unforgiving systems, with "invisible" and/or "untouchable" boundaries, may lead to systematic erosion of safety margins or to episodes where actors inadvertently break the envelope of safe operations. Opaque systems where safety-critical decisions are highly distributed may lead to situations where some actors influence the operational boundaries of other actors in ways that are not noticed by the latter.

8.3 Critical assumptions

In order to focus on the contrasts among the perspectives, we have tried to identify critical assumptions underlying each perspective (Table 5).

The energy and barrier perspective seems to entail two critical assumptions. The first one is that all accidents involve energy flows out of control as the immediate cause of harm. The second assumption is that effective risk control strategies can be derived from this energy view. We discussed these assumptions in Section 3.6.

Perrow's (1984) Normal Accident Theory makes three assumptions related to structural properties of socio- technical systems: (1) Systems with high interactive complexity can only be effectively controlled by a decentralised organisation. (2) Systems with tight couplings can only be effectively controlled by centralised organisations. (3) An organisation cannot be centralised and decentralised at the same time.

The following three claims are central to HRO Theory: (1) It is not feasible to totally eliminate erroneous actions. (2) However, it is possible to achieve nearly faultless performance by developing organisational redundancy, so that nearly all-erroneous actions are recovered before lead to severe harm. (3) Some organisations adapt to extreme demands by changing to an informal, competence based organisational structure and by adopting an informal interaction style.

The basic assumption of Turner's theory of man-made disasters is that there is a close link between the physical events involved in major accidents and the way the organisation handles information on hazards.

The perspective on conflicting objectives, as presented here, entails one central assumption: Human activities tend to migrate toward the boundary of acceptable performance as the actors search viable trade-offs between such considerations as workload and profitability.

At this level, the perspectives clearly point in different directions. Any perspective does not explicitly contradict the idea that accidents involve uncontrolled flows of energy. However, the other perspectives tend to entail a claim that organisational structure or processes must be included if we are to give a fruitful account of major accidents. Some researchers (e.g., Sagan,

1993) view Normal Accident theory and HRO theory as contradictory, whereas others (e.g. Rasmussen, 1994a) view them as complementary.

8.4 The relationship between major and minor accidents: The popular version of the iceberg theory

Few principles have been cited more often by safety practitioners than the iceberg theory. This principle was proposed by Heinrich (1931)[26]. For a given activity he compared the ratios between:

1. The number of insurance claim injuries (e.g. workers crossing the rail tracks to get to work between rail wagons),
2. Minor accidents from different scenarios related to the same activity (e.g. workers tripping over rails or stumbling on the uneven ground), and
3. The number of opportunities for an accident to occur (e.g. the number of times people crossed tracks). The "no injury" category was thus a measure of exposure, and not a frequency of near misses.

Heinrich found that the ratios were very variable, depending on the activity. The published ratios shown in (Figure 12) should therefore be interpreted as averages across a broad range of activities.

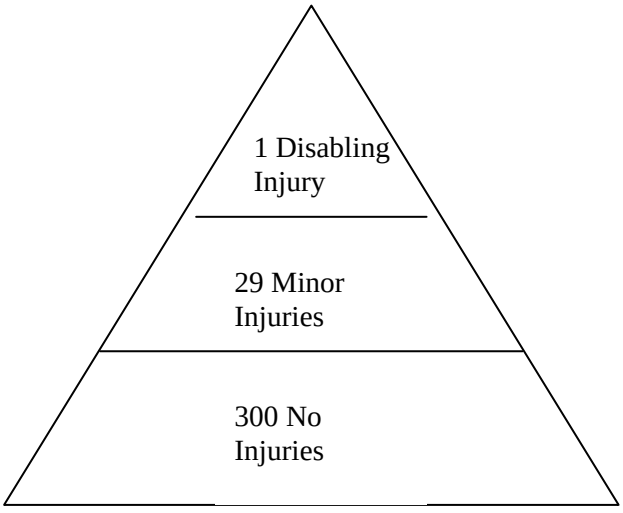

Figure 12. The Iceberg theory. Adapted from Heinrich (1931), cited in Hale (2000).

The popular understanding is that major accidents, minor accidents, and no-injury accidents arise from the same causal factors, i.e. unsafe practices and unsafe conditions. Risk reducing measures directed at these causal factors would influence major accidents, minor accidents and no-injury accidents to the same degree. Heinrich did *not* claim that the underlying causes for each degree of seriousness were the same.

[26] Our presentation of Heinrich's original version of the iceberg model draws heavily on a paper by Andrew Hale (2000).

The original statements of the iceberg theory thus made no strong claims concerning the relationships between major and minor accidents. However, a rather different version of the iceberg model has become strongly established among safety practitioners. According to this popular version, near misses, minor accidents and major accidents (including those with many fatalities) stem from the same causes. Moreover, according to this popular version, the ratios between major accidents, minor accidents and near accidents are constant. This implies that a preventive measure that successfully reduces the frequency of minor accidents will reduce the frequency of major accidents by a similar proportion. Another implication of the popular version is that low LTI[27] rates can be taken as a trustworthy indicator of the risk of major accidents.

Hale (2000) bluntly labelled the popular version of the iceberg theory an urban myth. He concluded that the available empirical evidence strongly indicates that major and minor injuries differ systematically with regard to locations, activities, amounts of energy released and the number and types of barriers or recovery opportunities that were bypassed in the event.

The perspectives outlined in this report may be used to further illuminate why the popular version of the iceberg theory is not plausible when it is extended to major accidents involving several fatalities. This amounts to asking what the perspectives have to say about the similarities and differences between major and minor accidents.

According to the energy and barrier perspective, minor and major accidents are both caused by uncontrolled energy releases. However, major accidents usually involve larger amounts of energy and the failure of multiple barrier functions, whereas minor accidents may arise from a single failure. Major accidents may also be related to unexpected dependencies among barriers.

In Normal Accident theory, this contrast is captured by the distinction between component failure accidents and system accidents. Normal Accident theory goes one step further, by relating system accidents to structural problems – either a highly interactive system controlled by a centralised organisation, or a tightly coupled system controlled by a decentralised organisation.

We are not aware of explicit discussions of this issue among HRO theorists. Their position seems to imply that a major accident in a HRO would typically involve a failure of organisational redundancy, i.e. of the mechanisms that usually recover errors before they lead to catastrophic outcome. A minor accident might occur because recovery mechanisms were not established in the first place.

We are not aware of explicit discussions of this issue in Turner's work on Man-made disasters either, since his discussions concentrate on major accidents. He implies that many smaller accidents can be indicators for disasters. He does not seem to imply that minor accidents arise from a long-term process of deteriorating information processing. Major accidents tend to come as fundamental surprises (Woods, 1990), whereas most minor accidents resemble earlier accidents experienced by the organisation. Prevention of minor accidents can thus be based on epidemiological approaches, whereas the prevention of major accidents requires analytic or evolutionary strategies (Rasmussen, 1997).

From the perspective of conflicting objectives, minor and major accidents are both related to the tendency of human activities (including decision-making) to migrate towards the boundary of acceptable performance (Section 7.2). However, the adaptation processes leading to migration may take place under different conditions (Rasmussen, 1994b). In manual activities, which are often associated with minor accidents, adaptation often takes place in a feedback loop with a

[27] LTI – Lost Time Incidents.

single operator and his equipment, tools or materials and their environment (e.g. a car driver adapting the speed to the environment, based on cues such as speedometer readings, speed limit signs, the quality of the road, traffic density and feedback from the vehicle). Organisational accidents in distributed systems typically involve several actors, each seeking local optimisation based on an incomplete view of the system. Major accidents thus tend to arise from situations where separated adaptation processes interact in a way that was not foreseen by the actors.

Taken together, these perspectives present several reasons why one should not expect major and minor accidents to arise from very similar patterns of safety problems. A company that displays excellent LTI records could be heading for a disaster, since some of the factors that contribute to major accident risk have little impact on the LTI frequency.[28] As a consequence, we need to consider the implications of these perspectives for prevention of major accidents. This is the topic of the next chapter.

[28] See Hopkins (2000b) for an example.

9 From theory to practice: Implications for risk control and accident prevention

We concluded the previous chapter with a claim that preventive strategies that work well with minor accidents may fail to provide effective control of major hazards. In this chapter we will summarise practical implications of the perspectives presented the previous chapters: How can we monitor major hazards? What risk reduction strategies can we derive from the perspectives? How can we learn the right lessons from disasters, minor accidents and near misses? What are the impacts of organisational change on major accident risks?

9.1 Monitoring the risk of organisational accidents

Effective safety management in complex and dynamic organisations is inconceivable without adequate feedback mechanisms (Kjellén, 2000:114). This is not only a matter of providing managers with appropriate information to support rational decision-making. Persistent feedback helps attract the attention of managers and employees to safety. Compelling feedback can also be necessary to "unfreeze" a rigid and unrealistic set of shared assumptions associated with the incubation period preceding a disaster (see Section 6.1). In this section we will discuss how the risk of organisational accidents can be monitored, with a view to the five theoretical perspectives.

The most straightforward way to monitor the risk of *minor* accidents in a stable system is to use loss-based performance indicators such as Lost Time Incident rates or Severity rates[29]. However, we should not confuse monitoring of risk with monitoring of loss. When speaking about risk, we usually refer to a *potential* for unwanted events in the future.[30] In contrast, the loss-based performance indicators express experienced loss in the past. Still, given a reliable reporting system, it makes good sense to interpret a statistically significant trend in experienced loss related to minor accidents as an indication of a corresponding change in the risk of minor accidents. Even a non-significant deterioration in a loss-based performance indicator may lead us to take action if we consider it better to react on a "false alarm" than to wait for enough aversive data to confirm a significant trend.

It is less straightforward to monitor the risks of major accidents. Major accidents are rare events, even if we consider highly aggregated data (e.g. "all accidents with more than five fatalities in the Norwegian petroleum industry"). The time needed to detect a significant trend in major accident frequency is generally far too long to provide useful feedback in a safety management context. Major accidents are also too rare to attract continuous management attention.

A common answer to this problem is to devise safety performance indicators based on information about contributing factors and root causes (Kjellén, 2000:248ff). One approach is to rate the elements of a company's safety management systems with reference to a model of an ideal safety management system. This approach is used by the *International Safety Rating System* (Bird and Germain, 1990), the Safety Element Method (Alteren, 1999) and Tripod Delta (van der

[29] The Lost-time injury rate is the number of lost time incidents per 106 employee-hours. A lost-time injury is an injury due to an accident at work, where the injured person does not return to work on the next shift. The Severity rate (S-rate) is the number of working days lost due to lost-time injuries per 106 employee-hours. Fatalities and 100 per cent permanent disability account for 7500 days (Kjellén, 2000).

[30] A thorough discussion of the risk concept is beyond the scope of this report, but it may be worth noting that risk may be viewed in two fundamentally different ways: (1) as a property of the world, which exists independently of the person who makes statements about the risk, and (2) as an expression of uncertainty of observable quantities of the world. See Aven (2003) for a discussion advocating the latter position.

Want, 1997). Can the perspectives presented in this report help us identify relevant indicators or conditions that could be followed up in order to identify changes in the risk of organisational accidents?

The *energy and barrier perspective* emphasises the fact that barrier functions need to be monitored and maintained. Barrier elements can deteriorate and interdependencies between barriers may arise or increase. Risk monitoring could be performed through investigating attributes of each individual barrier function, and of the relationship between them:

- How can a specific barrier function be taken care of? What are the important barrier elements?
- How can barrier functions fail, and are there interdependencies with other barrier functions (common cause failures)?
- How can barrier functions deteriorate?
- How can barrier functions be maintained and monitored?
- Are there potential indicators to measure availability/ efficiency?

A major challenge in monitoring barrier functions is to handle the diversity of different ways a barrier function may be carried out, and the great number of barrier elements involved. A quantitative approach for risk monitoring is the use of risk indicators. These may be based on the causal models in Quantitative Risk Analyses (Øien, 2001). The total set of risk indicators (technical and organisational) can be used to estimate the total change in risk, in the time periods between updates of the QRA.

Table 6. Practical implications of the perspectives on resilient organisations.

Issue	Energy and barrier perspective	Normal accident perspective	HRO perspective	Information processing perspective	Conflicting objectives, adaptation and drift
How can major accident risk be monitored?	Monitor the quality / effectiveness of barrier functions	Monitor interactive complexity and tightness of coupling. Monitor compatibility between control structure and technology.	Monitor the structural and cultural preconditions for organisational redundancy.	Combine HazOps and holistic approaches that include human and organisational factors. Check the organisation's ability to follow up signs of danger, e.g. near-accidents.	Is the feedback between different levels (vertical) from government to the daily immediate and open? Measurements/ strategic tools like balanced scorecard taking account of risk can be of help.
Risk reduction strategies	Include barrier functions in the design of the system; strive for 'defence-in-depth'. Ensure that compensatory measures are taken when barriers are unavailable. Monitor and maintain barrier functions throughout system life.	Reduce complexity or "loose" couplings. Apply decentralised control in systems with high interactive complexity, centralised control in tightly coupled systems. Discard high-risk systems that are both complex and tightly coupled.	Build organisational redundancy Build cultures that combines requirement for fault-free performance with openness to the fact that errors do occur.	Make systematic efforts to collect and analyse information about hazards and keep important hazards on the agenda. Build a culture that promotes active search for signals of danger, and knowledge sharing across organisational boundaries.	Make boundaries to unacceptable performance visible and touchable. Train personnel in boundary handling. Establish "counter-pressures" that favour safe performance. Keep a focus on situations where each stakeholder has a limited overview over the overall situation.
How can we learn from disasters and incidents?	We will see the effects of barriers and the barriers can be improved. Incidents can inform us about unexpected interactions or dependencies between barriers.	Incidents can provide information on unexpected interaction. However, single failure incidents provide little information on the structural problems that make a system prone to have system accidents.	Primarily, we learn from the daily operation and the normal procedure, but incidents/ accidents may demonstrate the absence of structural or cultural preconditions for organisational redundancy.	Arguing for use of several organisational learning approaches (e.g. Argyris & Schön [1978] double-loop learning where procedures for gathering signals about hazards are directly challenged). Focus on social learning, rather than individual learning.	Causal paths in an incident should be tracked beyond operator errors back to the normal operations in organisational units that contributed to create the incident scenario. Based on several incidents, one may create a work support system that makes decision-makers aware of the potential side effects.

Issue	Energy and barrier perspective	Normal accident perspective	HRO perspective	Information processing perspective	Conflicting objectives, adaptation and drift
How can technological and organisational change influence risk levels?	Increase of scale or speed may lead to increased accumulations of hazard sources. Maintenance and monitoring of barriers may deteriorate if requisite resources are no longer provided.	The degree of interactive complexity and coupling can change the preconditions for controlling the technology. Organisational change can cause incompatibility between control structure and technology.	Downsizing may affect the preconditions for organisational redundancy. HROs are capable of spontaneous reconfiguration in response to high demand or crisis.	Organisational change can make information flow more complex and indirectly influence risk. A problem solving process may be terminated because the relevant decision forum disappears.	Vertical communication lines can be weakened or disappear. Boundaries can be exceeded individually or by several actors at the same time. The systems of management and regulation may be unable to keep up with a fast pace of change in technology.

According to *Normal Accident theory* (Perrow, 1984), the organisation should ensure that the control structures are compatible with the properties of the technology. For instance, in order to cope with a tightly coupled technology, an organisation needs to maintain a correspondingly tight control structure. This may prove a challenging task if tight controls are not instrumental to carrying out the primary tasks (e.g. production). In this case, performance is likely to drift away from prescriptions given in rules and regulations as local adaptations take place. Moreover, the centralised control structures that are needed to handle emergencies may fail to materialise if they are not prepared and reinforced, e.g. through emergency drills. On the other hand, an organisation that handles technology characterised by complex interactions may need to monitor its capacity to improvise in a safe and effective manner in case of disturbances in the production process. There is also a need to monitor the technical modifications, since such modifications may introduce increased complexity or tighter couplings.

Monitoring of major accident risks following the *HRO theory* might focus on organisational redundancy in safety critical tasks: (i) Are the structural pre-conditions of organisational redundancy present? (ii) Does the organisational culture contribute to organisational redundancy? These issues are of current interest in connection with downsizing of organisations, where the structural condition of organisational redundancy is vulnerable, due to the fact that fewer persons are present, and at the risk of less overlap in knowledge. Are the preconditions for adequate handling of normal operations, as well as handling of deviations and emergency situations present in the organisation? The cultural dimension of organisational redundancy is not easy to monitor directly. A possible approach might be to analyse near misses, and accomplish a working environment survey, focusing on questions about loyalty, solidarity, frankness, authority gradient and communication.

HRO theorists also draw attention to the organisation's capacity to adapt to unforeseen situations. In this context Weick and Sutcliffe (2001) have devised a set of simple questionnaires devised to assess the degree of "mindfulness" in an organisation.

Risk monitoring, based on the *Information processing perspective* of accidents, may combine HazOps with holistic approaches that include human and organisational factors (Turner and Pidgeon, 1997). Turner and Pidgeon (1997: 187-189) emphasise the role of safety culture as a key for handling and continuously monitoring risk. They propose that the following issues should be focused in risk monitoring: i) senior management commitment to safety; ii) shared care and concern for hazards and their impact upon people; iii) realistic and flexible norms and rules about hazards; and iv) continual reflection upon practice through monitoring, analysis and feedback systems. A possible objective for monitoring efforts could be to assess the organisation's responses to signs of trouble, e.g.

- How does the organisation respond to concerns and warnings from outsiders (e.g. clients, contractors, media, regulatory authorities, NGOs)?
- What measures are decided and implemented in response to incidents and accidents? (None? Cosmetic fixes? Blaming the victim? Change of work practice or hardware? Change in managerial practices?)
- How are people voicing concern ("whistleblowers") treated? (Ignored? Ostracised? Celebrated?)

The perspective *Conflicting objectives, adaptation and drift* focuses on the capacity to handle conflicting objectives without drifting into dangerous system states. Uncontrolled local adaptations may lead to catastrophic and fundamentally surprising events in systems characterised

by distributed decision-making. According to Rasmussen's "migration" model, the boundaries of safe performance should be made visible and touchable, so that the actors have a way to know when they approach or exceed the boundaries. In situations with conflicting demands, people may face incentives for taking short cuts. Rasmussen also emphasises the need for feedback between different levels, from governments to the shop floor. In addition to the vertical dimension, we may be faced with cooperation between different companies at the operational level (following outsourcing, use of contract work). This will demand monitoring directed at cross border activities.

Risk monitoring, following this perspective, should emphasise the following questions: Is it possible for a single actor to know if he or she exceeds the boundaries for safe behaviour (does the system give distinct responses)? Will he/she be able to recover in case of exceeding the boundaries (or is the environment 'unforgiving')? Is the information level (e.g. the number of alarms) adapted to the human capacity? Other questions for review are about the relationship between safety and economy. Do managers follow up safety performance to the same extent as economy? Are high-level decision-makers held accountable for the accident risks associated with their decisions? Does risk monitoring include cross border operations when more than one enterprise is present (e.g. in the case of outsourcing)?

We should realise that performance indicators are two-edged swords. A negative trend can alert the organisation to problems, whereas a positive trend can lead managers to conclude that safety has been taken care of, and that they can direct their attention to other issues. Reliance on a single indicator can lead managers to focus on one aspect of health and safety (e.g. lost time incidents) and at the same time pay less attention to other health and safety issues (e.g. potential for major accidents; see Hopkins, 2000b for an example). In general, the persons whose performance is monitored tend to adapt to their behaviour to the performance indicators that are used.

9.2 Risk reduction strategies

An important advantage of applying more than one perspective on organisational resilience is that each perspective contributes to possible risk reduction strategies. By combining several perspectives, we can thus build a larger repertoire of risk reduction strategies. To gain optimal effect from risk reduction measures we have to ensure that structural and cultural aspects of the risk reduction strategies are compatible. The structural basis for safety control may consist of rules, regulations and working procedures, authority and responsibility assignment, reporting systems, formal communication, risk assessments and routines for deviation control. Cultural aspects of safe work performance include activities like ensuring employee involvement, knowledge sharing and organisational learning.

The *energy and barrier perspective* focuses on limiting energy amounts and controlling energy flows, as illustrated by Haddon's ten strategies for loss reduction (Section 3.1). Barrier functions are designed into technical systems and operational procedures. Administrative barriers can be seen as a part of this tradition. One example is the work permit system. No single barrier is 100 % effective, because they may have weaknesses due to active failures or latent conditions. Therefore, one way to build more safety into a technical system is to introduce multiple barriers, or a 'defence-in-depth' strategy (Reason, 1997). To be effective, this approach requires that interdependencies between barrier functions are minimised. The approach also requires provisions for monitoring and maintaining barriers.

According to *Normal Accident theory*, the preferred risk reduction strategy is to modify the technology in ways that reduce interactive complexity and loosen the couplings. This is the only effective strategy with systems that are both highly interactive and tightly coupled. An alternative

strategy is to adapt the organisation to the technology: Decentralise control with high interactive complexity, centralise with tight couplings. In this context, decentralisation implies more than just changing formal authorities. Local agents must be given the resources (e.g. information, competence, manpower) that are necessary to cope with the situations they may face.

The *HRO perspective* focuses on being proactive and to predict and prevent potential dangers as early as possible. A central risk reduction strategy is to build organisational redundancy. This strategy requires that a sufficient number of competent personnel are available, so that some overlap in competence, responsibilities and possibilities for observation is achieved. Workplace design should allow, and even encourage, seeking counsel from a colleague, observation of other people's work, and intervention in case of an erroneous action. Moreover, it is necessary to build a culture that encourages questioning and intervention. Another strategy is to build organisations with a capacity for spontaneous, adaptive reconfiguration.

The *information processing perspective* puts a strong focus on the gathering, interpretation and dissemination of information. One aspect of this is the systematic collection, analysis and dissemination of information about hazards, and actively trying to find out what we do not know. Another aspect is the building of "cultures with requisite imagination", i.e. cultures which encourage sharing of information, innovation and learning (Westrum, 1993).

With regard to *conflicting objectives*, three risk reduction strategies can be derived from the migration model (Section 7.2, p. 44): The first is to make boundaries towards unacceptable risk visible to the relevant actors. This can be challenging in practice, since we tend to adapt to many kinds of warnings, so that the boundary can become "invisible" over time. The second strategy is to make boundaries touchable. This implies that the actor is given a chance to recover if he strays beyond the boundary to unacceptable performance. The third strategy is to provide a "counter-pressure" that favours safe actions, e.g. through follow-up and feedback. In situations with *distributed decision making,* when a lot of activity is taking place and each stakeholder has a limited overview, there is a possibility that decisions perceived as safe by each local actor may interact in unforeseen ways and trigger an accident. Such situations may in principle be avoided by giving the actors access to more information, thus improving their situation awareness, provided that the actors are in a position to handle the increased information load. Alternatively, one may try to program the decisions to be made, i.e. specify in advance the action to be taken by each actor so as to make their actions predictable. The latter approach is characteristic of railway operations. Many actors (e.g. train drivers) have a very restricted view of the total system, but their actions are tightly controlled through regulations, orders and signalling systems so as to avoid conflicts.

The responsibility for safety must be clearly communicated. If a strategic decision is perceived to cause trouble, in let us say the maintenance department, this uncertainty and perception must be immediately communicated back to the management level. This will make it easier to see the relationship between decisions/ safety measures and the risk level, and will also serve as a tool to establish counter pressure.

The adaptation perspective also alerts us to the possibility that some risk reducing measures may lead to behavioural change ("compensation") that reduces the risk-reducing effect (Wilde, 1982). In extreme cases, a person may overestimate the effect of a risk-reducing measure and change his/her behaviour to a point where the risk is greater than it was before the risk-reducing measure was introduced.

The risk reduction strategies derived from the different perspectives tend to be complementary, with some overlaps. There is, however, one noteworthy tension. Uncritical application of the

energy and barrier perspective might lead us to combine numerous physical, technical and organisational barriers in order to contain a hazard source. According to Normal Accident theory, this strategy could fool us into designing a system with high interactive complexity. A safety system may under adverse circumstances camouflage the source of a disturbance and cause operators to diagnose a problem incorrectly.

In some cases, a search for risk-reducing measures is triggered because a quantitative risk analysis (QRA) produces results that are not compatible with the acceptance criteria of the activity. In such situations, the search for risk-reducing measures is often restricted to factors that are explicitly modelled in the QRA. For instance, if an unacceptable risk level is identified in the QRA of an offshore installation, a common response is to propose a modification of the design and recalculate the risk, using the same QRA approach. Measures to reduce the frequency of hydrocarbon leaks in the first place are often not considered, because such measures will usually not have any impact on the calculated risk if standard QRA procedures are used.

9.3 Learning from disasters and precursors

The perspectives contribute concepts and directions that may help us to describe and explain accidents and incidents. They may also contribute some insight into the preconditions for learning from accidents.

The *Energy and barrier perspective* focuses on effects of barriers and examines how barriers could be improved, or if new barriers should be established. From this perspective, the organisation should use information on disasters and precursors to identify weak points and problems related to barrier functions. It should also use such information to identify and improve problems related to the monitoring and maintenance of barrier functions. Organisations should, for instance, pay close attention to incidents related to their work permit systems. Work permit systems can be safety critical because they often replace several technical barriers. Finally, one may ask whether the hazardous energy source is really necessary, whether the amount of energy could be reduced, and whether the hazard source is adequately controlled.

The *Normal Accident perspective* leads us to ask whether the presence of interactive complexity or tight couplings contributed to the accident. Incidents could also give the organisation information on how well its control structures are adapted to the properties of its technology. For instance, in the case of an organisation handling a tightly coupled technology, an incident or accidents may reveal gaps in the control structure. Information from incident investigations can thus be used to reduce complexity, loosen couplings, or at least to prepare the organisation to handle the problems created by complexity and tight coupling.

The *HRO perspective* covers both structural and cultural aspects of the organisation. LaPorte and Consolini (1991) stressed that normal situations and precursors, and not only disasters, are important sources for learning. Weick (1987) showed that some HROs (nuclear submarines) use story-telling to build a resilient culture. Stories about accidents and critical incidents have an important place in this socialisation process.

The *information processing perspective* invites us to look for deficiencies in the organisation's information handling prior to the incident or accident. Did someone have the pieces of information that were necessary to foresee and avert the accident? Why was this information not acted on? In some cases, the information that is needed to identify and assess a sign of danger is distributed among several disciplines or organisational units. The focus after accidents must be on social learning, rather than individual learning. People in the organisation should be trained to ask

critical question and investigate assumptions in what is termed as “public testing” by Argyris and Schön (1978).

The *conflicting objectives perspective* proposes that systems may drift towards greater risk as performance and decision processes are shaped by competitive pressure. Decision-makers in distributed systems may be running risks because they are not aware of how their decisions interact with the decisions made by other actors. Investigators of accidents and incidents should pay attention to human-system-interaction at the boundary of acceptable performance. They should also examine the way decisions are distributed among decision-makers and decision fora, looking for cases where decisions which interact strongly in their impact on risk are taken in different fora. Rasmussen and Svedung (2000) argue that causal paths in an incident should be tracked beyond operator errors and hardware failures, back to the normal operations in the organisational units that contributed to create the accident scenario. Based on several incidents one may create a work support system that makes decision-makers aware of potential side effects.

Accident investigators need to apply several perspectives in order to raise critical questions, since there is always a danger that you will find only what you are looking for. Important questions are:

- Who is represented in the accident investigation group? Does the composition of the group provide for knowledge sharing and the application of multiple perspectives?
- How do experts and lay persons understand the accident?
- What perspectives are used and why?
- Are we able to draw generic lessons from the specifics of a particular incident?
- How can we improve our ability to learn from incidents (second order learning)?

Two common pitfalls should be kept in mind. Organisations often focus too much on the idiosyncratic or atypical aspects of an accident. This may lead to measures that only apply to a specific activity in a specific place under specific circumstances, whereas similar problems in similar situations remain uncorrected. The other pitfall is to tighten procedures and rules to a point where they are not compatible with efficient performance of work. Unrealistic procedures and rules are not only ineffective. They also tend to produce discrepancies between word and action, and to foster a culture where such discrepancies are tacitly accepted.

9.4 Resilience and change

We noted in the introduction that risk management efforts currently face a very fast pace of change in technology, which is associated with increasing scales of industrial installations and high degrees of integration and coupling of systems. At the same time, organisational structures change dramatically in response to a very aggressive and competitive environment. Finally, the political and regulatory environment changes in various directions; sometimes leading to more elaborate demands for risk control (e.g. the Seveso directive). In this section we will consider organisational and technical change that has been initiated in order to achieve other goal than safety improvements.

The effects of change are obviously complex. At the organisational level, one needs to distinguish between the immediate effect of the change process as such (e.g. employee anxiety about loosing their jobs), and the long-term effect of the changed characteristics of the organisation (e.g., lower manning levels, new control structures). Pfeffer (1998) argued that, while the short terms effect of a downsizing in saved costs are quite easy to achieve and measure, the long-term effects are more uncertain. Lee Marks (1993) has surveyed the consequences of downsizing in US industry, see Table 7:

Table 7. Effects of downsizing (Lee Marks, 1993).

Consequence	Percent of firms reporting
Lower morale among remaining work force	61
More need for retraining remaining employees	41
More use of retraining remaining employees	36
More use of temporary workers and contractors	35
Increased retiree health care costs	30
Entire functions contracted out	25
Wrong people lost	20
Severance costs greater than anticipated	16
Too many people lost	6

These numbers are averages across many industries, and thus hide variations. However, they demonstrate some of the typical effects on downsizing. The high percentages of firms reporting lower morale and retraining needs are noteworthy.

How can technological and organisational change affect the resilience of organisations? From the *energy and barrier perspective* we may note that increases of scale or speed may be associated with increased accumulations of hazard sources (e.g. vehicles that carry more kinetic energy; plants with larger repositories of dangerous substances). However, change may also go in the opposite direction. Efforts to reduce consumption of input factors may, e.g., lead to intensification of processes so that smaller amounts of dangerous substances need to be handled (Kletz, 1991). Inventories of dangerous substances and goods may be reduced to reduce inventory costs. The number of people exposed to a hazard may be reduced as a consequence of reduced manning levels.[31] Change can also affect the resources needed to effectively monitor and maintain barriers, such as competent manpower. From this perspective, one may want to monitor the backlog on preventive maintenance of safety barriers.

Normal Accident Theory directs attention to changes in technology as well as control structure. Technological change may affect the degree of interactive complexity or lead to tighter coupling. Organisational change may lead to either more or less centralised control structures. As a consequence, the degree of compatibility between technology and control structure may change. Change sometimes leads to reduced risk levels. For instance, Perrow (1984: 159ff) argues that safety in the U.S. system of airways (including Air traffic control) was improved during 1960s and 1970s due to reduced coupling and complexity. However, it seems plausible that the dominant trends are towards tighter coupling and more interactive complexity. Indeed, Beck (1992, 1999) claims that such changes are part of major transformation of society ("reflexive modernisation"). New risks, such as those related to the greenhouse effect and the nuclear power industry, are global and can strike everyone.

Many of the organisations studied by researchers in the *HRO* tradition are stable and conservative (e.g. Bierly, 1995). This research tradition thus provides limited empirical evidence on the effects of change. However, HRO theory directs attention to possible effects of downsizing on

[31] The impacts of reduced manning level on risk can be complex. The expected number of injuries and fatalities associated with the activity may decrease because fewer persons are exposed to the hazard. The individual risks of the remaining workers may remain unchanged, or perhaps increase somewhat if fewer persons have to perform the same set of hazardous tasks. The over all effect for society depends on what happens to the persons that are made redundant – e.g. whether they enter new jobs that are less risky. These are only the impacts of changes in exposure. The over all impact will also depend on whether risk control improves or deteriorates as a result the organisational change.

organisational redundancy (Rosness et al., 2000). The instrumental preconditions for redundancy may, e.g., be threatened as a result of reduced manning levels. Going from two control room operators to one may imply that nobody may be there to intervene in case the remaining control room operator responds inadequately to an alarm. Outsourcing may lead to cultural barriers in the organisations, and in some cases weaken the cultural preconditions for organisational redundancy. Organisational change may, on the other hand, be associated with a reduction of authority gradients and thus facilitate communication and cooperation.

The *information processing* perspective emphasises impacts of technological change on the flow of safety-relevant information. This includes the quality of communication channels – the impact of face to face communication is sometimes very different from that of an e-mail or a memo (Weick, 1987). Adoption of information and communication technology thus may have profound impacts on organisational resilience. Organisational change may also affect ongoing problem solving processes. For instance, a decision forum may disappear during a change process, and thus leave some safety issues unresolved. Moreover, organisational change may affect the resources and status of personnel serving as "watchdogs" representing safety interests. These tasks are often associated with staff functions, and may be seen as "non-productive" in organisations undergoing a business process reengineering process.

The *conflicting objectives* perspective emphasises the potential for drift: Many small changes may accumulate and gradually cause a system to migrate beyond the boundary of safe operations. Decision processes that interact strongly in their impact on safety may be allocated to different decision arenas, leading to the possibility of conflicting decision outputs. The pace of technological and organisational change may exceed the speed of information handling in regulatory and safety management processes. A high pace of technological change renders retroactive control strategies ineffective, since experience becomes obsolete at a correspondingly high pace. A common regulatory strategy is to replace detailed prescriptive regulation with goal-oriented legislation (Hopkins and Hale, 2002). This strategy allows enterprises to develop and implement new means to accomplish given safety objectives without waiting for new legislation to approve the new solutions.

Based on experience from the Norwegian oil industry, Serck-Hansen and Steinum (2001) claim that the quality of the change process is essential for its safety outcome. They emphasise the quality of participation, founded on trust and a common understanding of the situation. Focus should be kept on tasks in the design phase. They call for patience in implementation and warn against too early focus on the final results.

Taken together, the perspectives show the diversity and complexity of the possible impacts of organisational and technological change on safety. We have to realise that it is not feasible to predict all safety impacts of a major change process. This situation calls for robust safety management strategies, where proactive evaluations of the proposed change process and its proposed outcome are combined with continuous monitoring of the effects of the process. For instance, a normative safety management model such as the SMORT questionnaire (Safety Management and Review Technique; Kjellén et al., 1987; Kjellén, 2000) may be used upfront to detect instances where the proposed organisation fails to provide the resources and routines that are needed to take care of safety. SMORT also contains checklists pertaining to the planning and execution of projects. During and after the process, interviews and questionnaires may be used to monitor its effects. Moreover, incident and accident investigations may focus on the impacts of organisational change on safety. In order to succeed, this approach requires top management commitment to safety. Top management must be willing to change their plans, moderate their ambition level, postpone the process and even reverse some of the modifications if this is necessary to achieve the company's safety objectives.

9.5 Epilogue

At this point, some readers may miss a tight and elegant synthesis of the diverse perspectives and ideas on organisational accidents and organisational resilience. Why not put everything into a neat little model, or at least compile a handy checklist, so that people can use the ideas without bothering about five different perspectives?

Our answer is that the five perspectives, with all their associated concepts and propositions, do not fit into a neat little model. The concepts and ideas of different perspectives cannot always be readily expressed using the concepts of a different perspective. Neither do we believe that it is simply a question of pitting the perspectives against each other, devising empirical tests, and deciding which perspective most adequately reflects the realities of organisational life[32].

Moreover, we do not believe that simplification and reduction is always a good thing in the life of organisations. This brings us back to Westrum's notion of requisite imagination and Schulman's argument in favour of conceptual slack (Section 6.2 above; see also Westrum, 1993; Weick, 1987; Schulman, 1993). We suggest that complex organisations in a dynamic, ambiguous environment thrive on open-minded controversies, rather than single-minded consensus. Open-minded organisations with room for divergent opinions may respond more effectively to signs of trouble than organisations with less tolerance for divergent interpretations of reality.

Finally, we contend that a rich repertoire of perspectives is a great asset to the safety practitioner. It is extremely difficult to get attention to the same message year after year. With a broad array of perspectives at his or her disposal, the safety practitioner is in a much better position to bring out new messages, to surprise and to provoke interest.

[32] The reader may have noted that Sagan (1993) tried to do something along these lines in book "Limits to safety". Although his results are interesting in their own right, they hardly provide a refutation of HRO theory; see Section 5.8 above and Rasmussen (1994a).

10 References

Alteren, B. (1999): Implementation and evaluation of the Safety Element Method at four mining sites. *Safety Science, 31*: 231-264.

Argyris, C. and D. A. Schön (1978): *Organizational Learning*. Reading, Massachusetts: Addison-Wesley.

Ashby, W.R. (1981): Self-regulation and requisite variety. In F.E. Emery (ed.): *Systems Thinking. Volume One.* Harmondsworth: Penguin Education, 100-120. Earlier published as Chapter 11 in W.R. Ashby (1956): *Introduction to Cybernetics,* Wiley.

Aven, T. (2003): *Foundations of Risk Analysis – A knowledge and decision oriented perspective.* Chichester: Wiley.

Bainbridge, L. (1987): Ironies of automation. In J. Rasmussen, K. Duncan and J. LePlat (eds.): *New Technology and Human Error*. Chichester: Wiley (271-283).

Beck, U. (1992): *Risk Society: Towards a New Modernity.* London: Sage.

Beck, U. (1999): *World Risk Society.* Cambridge: Polity Press.

Bierly, P.E. and Spender, J.-C. (1995): Culture and High Reliability Organizations: The case of the nuclear submarine. *Journal of Management, 21* (4), 693-656.

Bird, F.E. and Germain, G.L. (1985): *Practical Loss Control Leadership.* Institute Publishing, Division of International Loss Control Institute, Loganville, Georgia.

Bourrier, M. (1998): Elements for designing a self-correcting organisation: Examples from nuclear power plants. In A. Hale and M. Baram (eds.): *Safety Management. The Challenge of Change*. Oxford: Pergamon.

Brehmer, B. (1991): Distributed decision making: Some notes on the literature. I J. Rasmussen, B. Brehmer og J. Leplat (eds.): *Distributed decision making: Cognitive models for cooperative work.* Chichester: Wiley.

Burns, T. R. and G. M. Stalker (1961): *The Management of Innovations*. London: Tavistock.

Clarke, L. and Perrow, C. (1996): Prosaic Organizational Failure. *American Behavioral Scientist, 39* (8) 1040-1056.

Cullen, L. (1990): The Public Inquiry into the Piper Alpha Disaster. London: HSO.

Foster, H.D. (1993): Resilience theory and system evaluation. In J.A. Wise, V. D. Hopkin and P. Stager (eds): *Verification and Validation of Complex Systems: Human Factors Issues.* Berlin: Springer, 35-60.

Gibson, J. J. (1961): The contribution of experimental psychology to the formulation of the problem of safety – a brief for basic research. In *Behavioral Approaches to Accident Research,* New York: Association for the Aid of Crippled Children, pp. 77-89. Reprinted in W. Haddon, E.A. Suchman and D. Klein (1964): *Accident Research: Methods and Approaches.* New York: Harper & Row.

Guttormsen, G., Randmæl, S. and Rosness, R. (2003): *Utforming av regelverk for togframføring.* Report STF38 A03408. Trondheim: SINTEF Industrial Management. (Design and formulation of operational rules for railways. In Norwegian.)

Haddon, W. (1970): On the escape of tigers: An ecological note. *Technological review, 72* (7), Massachusetts Institute of Technology, May 1970.

Haddon, W. (1980): The Basic Strategies for Reducing Damage from Hazards of All Kinds. *Hazard prevention,* Sept./ Oct. 1980.

Hale, A. (2000): Conditions of occurrence of major and minor accidents. 2me séance du séminaire "Le risque de défaillance et son contrõle par les individus et les organisations", 6-7 novembre, Gif sur Yvette.

Heinrich, H. W. (1931): *Industrial Accident Prevention.* New York: McGraw-Hill. (Cited by Hale, 2000).

Hollnagel, E. (1999): *Accident analysis and barrier functions.* Halden, Norway: Institute for Energy Technology.

Hopkins, A. (1999): The limits of Normal Accident theory. *Safety Science, 32* (2-3), 93-102.

Hopkins, A. (2000a): An AcciMap of the Esso Australia Gas Plant Explosion. Paper presented at the 18th ESReDa seminar Risk Management and Human Reliability in Social context. Karlstad, Sweden, June 15- 16, 2000.

Hopkins, A. (2000b): *Lessons from Longford: The Esso Gas Plant Explosion.* Sydney: CCH. (Can be ordered by e-mail to salescentre@cch.com.au.)

Hopkins, A. and Hale, A. (2002): Issues in the regulation of safety: Setting the scene. In B. Kirwan, A. Hale and A. Hopkins (eds.): *Changing Regulation: Controlling Risks in Society.* Oxford: Pergamon.

Hovden, J. and Steiro T. (2000): The effects of cost- cutting in the Norwegian petroleum industry. Presented at *ESREL 2000, Foresight and Precaution*. In Cottam, Harvey, Pape & Tait (eds), Balkema, Rotterdam, ISBN 90 5809 140 6, P. 601-605

Johnson, W. G. (1980): *MORT Safety Assurance Systems.* New York: Marcel Dekker.

Kjellén, U. 2000: *Prevention of Accidents Through Experience Feedback.* Taylor & Francis, London and New York.

Kjellén, U., Tinmannsvik, R.K., Ulleberg, T., Olsen, P.E., Saxvik, B. (1987): *SMORT: Sikkerhetsanalyse av industriell organisasjon. Offshore-versjon.* [MORT. Safety analysis of industrial organisations. Offshore version.] Oslo: Yrkeslitteratur.

Kørte, J., Aven, T. and Rosness, R. (2002): On the use of risk analysis in different decision settings. Paper presented at ESREL 2002, Decision Making and Risk Management, Lyon, 19-21 March 2002.

Kletz, T. (1991): *Plant design for safety: A user friendly approach.* New York : Hemisphere Publ.

LaPorte, T. R. and Consolini, P.M. (1991): Working in practice but not in theory: Theoretical challenges of "High-Reliability Organisations". *Journal of Public Administration Research and Theory, 1*, 19-47.

Lee Marks, M. (1993): Restructuring and downsizing, In: Mirvis, P. H. (ed.): *Building the Competitive Workforce* (New York: John Wiley, 1993).

Lindblom, C. E. (1959): The science of "muddling through". *Public administration Review, 19*, 79-88.

March, J. G. and Olsen, J. P. (1976): *Ambiguity and Choice in Organizations.* Bergen: Universitetsforlaget.

NOU 2000: 30: *Åsta- ulykken, 4. januar 2000. Hovedrapport.* [The Åsta accident, January 4, 2000. Main report.] Justis- og politidepartementet. Statens forvaltningstjeneste, 2001. Electronic version available at http://odin.dep.no/jd/norsk/publ/utredninger/NOU/012001-020007/index-dok000-b-n-a.html

NOU 2001:9: *Lillestrøm-ulykken 5. april 2000.* [The Lillestrm accident, April 4, 2000]. Justis og politidepartementet. Statens forvaltningstjeneste, 2001.

Perrow, C. (1984): *Normal Accidents.* New York: Basic Books.

Perrow, C. (1986): *Complex Organizations. A Critical Essay*. New York: Random House.

Pfeffer, J. (1998): *The Human Equation: Building profits by putting people first.* Boston: Harvard Business School Press.

Pidgeon, N. and O'Leary, M. (2000): Man-made disasters: why technology and organizations (sometimes) fail. *Safety Science, 34*, 15- 30.

Rasmussen, J. (1986): *Information processing and human-machine interaction. An approach to cognitive engineering*. New York: North-Holland.

Rasmussen, J. (1994a): High Reliability Organizations, Normal Accidents, and other dimensions of a risk management problem. Paper. *NATO Advanced Research Workshop on Nuclear Arms Safety*. Oxford, UK, August 1994.

Rasmussen, J. (1994b): Risk management, adaptation, and design for safety. In B. Brehmer and N.-E. Sahlin (eds): *Future Risks and Risk Management*, (pp 1-36). Dordrecht: Kluwer Academic Publishers.

Rasmussen, J. (1996): Risk management in a dynamic society. Presentation at the seminar *Safety and Reliability in Industrial Management*, Trondheim 29-30 May 1996. (Viewgraphs)

Rasmussen, J. (1997): Risk management in a Dynamic Society:
A Modelling Problem *Safety Science, 27* (2-3), pp. 183-213.

Rasmussen, J. and I. Svedung (2000): *Proactive Risk Management in a Dynamic Society* (Swedish Rescue Services Agency, Karlstad, Sweden).

Reason, J. (1990): *Human error*. Cambridge: Cambridge University Press.

Reason, J. 1997: *Managing the Risks of Organizational Accidents*. Ashgate.

Rochlin, G. I., LaPorte, T. and Roberts, K. H. (1987): The self-designing high-reliability organization: Aircraft carrier flight operations at sea. *Naval War College Review 40*(4), 76-90. Also available on the Internet site:
http://www.nwc.navy.mil/press/review/1998/summer/art7su98.htm

Rosness, R. (2001): *"Om jeg hamrer eller hamres, like fullt så skal der jamres. Målkonflikter og sikkerhet."* [On conflicting goals and safety.] SINTEF Report STF38 A01408. Trondheim: SINTEF Industrial Management. Available at www.risikoforsk.no.

Rosness, R., Håkonsen, G., Steiro, T. and Tinmannsvik, R.K. (2000): The vulnerable robustness of High Reliability Organisations: A case study report from an offshore oil production platform. Paper presented at the 18th ESReDA seminar *Risk Management and Human Reliability in Social Context.* Karlstad, Sweden, June 15-16, 2000.

Sagan, S. D. (1993): *The limits to safety. Organizations, accidents, and nuclear weapons.* Princeton, New, Jersey: Princeton University Press.

Schulman, P. R. (1993): The negotiated order of organizational reliability. *Administration & Society, 25* (3), 353-372.

Serck-Hansen, C. and Steinum, T. (2001): Methodology for change: Results from a workshop on safety and organisational change. In E. Serck-Hansen (ed.): *Safe Change. Methodology on Change in Norwegian Oil Industry.* Høvik: Det Norske Veritas.

Statens jernbanetilsyn (1999): Forskrift om " Krav til styring og oppfølging....". [Regulations concerning "Requirements for control and follow-up…]

Svenson, O. (1991): The accident evolution and barrier function (AEB) model applied to incident analysis in the processing industries. *Risk Analysis, 11* (3) 499-507.

Turner, B. A. (1978): *Man-made disasters.* London: Wykeham Science Press.

Turner, B. A., Pidgeon, N. F. (1997): *Man-made disasters*. 2nd Edition. London: Butterworth-Heinemann.

Van der Graaf, G. (2001): Hearts and minds. Paper presented at *Sikkerhetsdagene 2001,* Trondheim 30-31 October 2001.

Van der Want, P.D.G. (1997): Tripod incident analysis methodology. In: J. van Steen (ed.): *Safety Perormance Measurement.* Warwickshire, UK: Institution of Chemical Engineers

Wagenaar, W. A. and Groeneweg, J. (1987): Accidents at sea: Multiple causes and impossible consequences. *International Journal of Man-Machine Studies, 27,* 587-598.

Wagenaar, W. A., Groeneweg, J., Hudson, P.T.W. and Reason, J.T. (1994): Promoting safety in the oil industry. *Ergonomics, 37,* 1999-2013.

Weick, K. E. (1987): Organizational culture as a source of high reliability. *California Management Review, 29,* (2) 112-127.

Weick, K. E. (1990): The vulnerable system: An analysis of the Tenerife air disaster. *Journal of Management, 16* (3), 571-593.

Weick, K.E. and Sutcliffe, K.M. (2001): *Managing the Unexpected.* San Francisco: Jossey-Bass.

Westrum, R. (1993): Cultures with Requisite Imagination. In J.A. Wise, V. D. Hopkin and P. Stager (eds): *Verification and Validation of Complex Systems: Human Factors Issues.* Berlin: Springer, 401-416.

Wilde, G.J.S. (1982): The theory of risk homeostasis: Implications for safety and health. *Risk Analysis, 2* (4): 209-225.

Woods, D. D. (1990): Risk and human performance: Measuring the potential for disaster. *Reliability Engineering and System Safety, 29,* 387- 405.

Woods, D. D., Johannesen, L.J., Cook, R.I., Sarter, N.B. (1994): *Behind Human Error: Cognitive Systems, Computers, and Hindsight.* State- of- the- Art Report 94- 01. Wright-Patterson Airforce Base, Ohio: CSERIAC

Yin, R. K. (1994): *Case Study Research. design and methods.* Second edition. Thousand Oaks: Sage.

Øien, K. (2001): *Risk control of offshore installations. A framework for the establishment of risk indicators.* Ph.D. Thesis. Department of Production and Quality Engineering. Trondheim: NTNU.