

EUROPEAN ORGANISATION
FOR THE SAFETY OF AIR NAVIGATION



EUROCONTROL EXPERIMENTAL CENTRE

REVISITING THE « SWISS CHEESE » MODEL OF ACCIDENTS

EEC Note No. 13/06

Project Safbuild

Issued: October 2006

The information contained in this document is the property of the EUROCONTROL Agency and no part should be reproduced in any form without the Agency's permission.
The views expressed herein do not necessarily reflect the official views or policy of the Agency.

REPORT DOCUMENTATION PAGE

Reference: EEC Note No. 13/06		Security Classification: Unclassified				
Originator: EEC - SAS (Safety Analysis and Scientific)		Originator (Corporate Author) Name/Location: EUROCONTROL Experimental Centre Centre de Bois des Bordes B.P.15 F - 91222 Brétigny-sur-Orge Cedex FRANCE Telephone : +33 (0)1 69 88 75 00Internet : www.eurocontrol.int				
Sponsor: EATM		Sponsor (Contract Authority) Name/Location: EUROCONTROL Agency Rue de la Fusée, 96 B -1130 BRUXELLES Telephone: +32 2 729 9011				
TITLE: <p align="center">REVISITING THE « SWISS CHEESE » MODEL OF ACCIDENTS</p>						
Authors J. Reason (for Dedale) E. Hollnagel (for Dedale) J Paries (Dedale)	Date 10/2006	Pages x+25	Figures 9	Tables --	Annexes 1	References 22
EEC Contact Fabrice Drogoul	Project Safbuild		Task No. Sponsor 120000-SRD-3-E1-0000	Period 2004-2005		
Distribution Statement: (a) Controlled by: Head of SRT (b) Special Limitations: None						
Descriptors (keywords): Swiss cheese model, accident model, Safety						
Abstract: <p>Accidents in complex system occur through the accumulation of multiple factors and failures. J. Reason has famously developed a model based on the Swiss Cheese Metaphor that suggests multiple contributors (the holes in cheese slices) must be aligned for any adverse events to occur. Barriers in a system (the slices themselves) are intended to prevent errors that result in these adverse events. This Swiss cheese model is not without drawbacks, and is not accepted uncritically. With use over time even the author has acknowledged its limitations. Nevertheless it remains widely used and is employed as the main basis for new method development. The concept of barrier provides one of the few opportunities to model interactions and complexity in High Risk domains. A compilation of pro's and con's was needed to gain insight into the potential and limitations of this model for ATM application.</p>						

FOREWORD

This report tries to fill this gap by showing the suitability and limits of this model.

The Swiss Cheese Model is heavily used in safety critical domain and in particular in ATM. Recently it has been subject to criticism but yet a review of the criticism and a discussion on pro's and con's of the model was not available.

This document gives the results of a short project to revisit the reason Swiss Cheese Model.

Fabrice Drogoul
Project Leader

Page intentionally left blank

TABLE OF CONTENTS

LIST OF FIGURES	VII
LIST OF ANNEXES	VII
REFERENCES	VIII
1. PREAMBLE	1
2. SCOPE AND PURPOSE OF THE STUDY	2
3. HISTORY OF THE MODEL	4
4. FUNCTIONS OF THE MODEL	9
4.1. SCM AS A CONCEPTUAL FRAMEWORK.....	9
4.2. SCM AS A MEANS OF COMMUNICATION	9
4.3. SCM AS A BASIS FOR ANALYSIS.....	10
5. DISSEMINATION OF THE MODEL	11
6. PUBLISHED CRITICISMS OF THE MODEL	12
7. FURTHER COMMENTS ON THE SWISS CHEESE MODEL (SCM)	14
7.1. THE STATUS OF THE SCM AS A MODEL	14
7.2. THE SCM AS A MEANS OF COMMUNICATION	14
7.3. NORMATIVE MODEL OF ORGANISATIONS	15
7.4. ORDERLINESS AS AN ARTEFACT OF RETROSPECTION	16
7.5. PROXIMAL AND DISTAL CAUSES	16
7.6. BARRIERS AND DEFENCES	17
7.7. CONTEMPORARY VIEW OF ACCIDENTS.....	17
7.8. THE SCM AS A BASIS FOR ANALYSIS	19
7.9. SCM AS A BASIS FOR MEASUREMENTS.....	20
8. CONCLUSIONS	21

LIST OF FIGURES

Figure 1	5
Figure 2	5
Figure 3	6
Figure 4	7
Figure 5: the Mark 3 version.....	8
Figure 6	10
Figure 7	10
Figure 8: SCM Mark I (pathological aspects).....	15
Figure 9: SCM Mark II	19

LIST OF ANNEXES

ANNEX A - Uberlingen accident analysis using the Reason model framework	25
---	----

REFERENCES

1. Bogner, M. S. (2002). Stretching the search for the “why” of error: the systems approach. *Journal of Clinical Engineering*, 27, 110-115.
2. Dekker, S. (2002). *The Field Guide to Human Error Investigations*. Ashgate.
3. Hannaman, G. W., Spurgin, A. J. & Lukic, Y. D. (1984). Human cognitive reliability model for PRA analysis (NUS-4531). Palo Alto, CA: Electric Power Research Institute.
4. Heinrich, H. W. (1931). *Industrial accident prevention*. McGraw-Hill.
5. Hollnagel, E. (1998). *Cognitive reliability and error analysis method*. Oxford: Elsevier Science Ltd.
6. Hollnagel, E. (2004). *Barriers and accident prevention*. Aldershot, UK: Ashgate.
7. Luxhoj & Kauffeld (2003). vol 5. *The Rutgers Scholar*, 2003).
8. Leveson, N. (2004). A New Accident Model for Engineering Safer Systems. *Safety Science*, Vol. 42, No. 4, April 2004, pp. 237-270.
9. Perrow, C. (1984). *Normal Accidents: Living With High-Risk Technologies*. New-York, USA: Basic Books, Inc.
10. Pidgeon, 1997. *Man-Made Disasters*. Second Edition. Butterworth-Heinemann.
11. Rasmussen, J., Pedersen, O. M., Mancini, G., Carnino, A., Griffon, M. & Gagnolet, P. (1981). Classification system for reporting events involving human malfunctions (Risø-M-2240, Reason, J. T. (1990). *Human Error*. Cambridge: Cambridge University Press.
12. Reason, J. (1990) The contribution of latent human failures to the breakdown of complex systems. *Philosophical Transactions of the Royal Society (London)*, series B. 327: 475-484.
13. Reason, J. T. (1997). *Managing the risks of organizational accidents*. Aldershot, UK: Ashgate Publishing Limited.
14. Reason, J., Shotton, R., Wagenaar, W.A., Hudson, P. T. W., & Groeneweg, J. (1989) *Tripod: A Principled Basis for Safer Operations*. The Hague: Shell Internationale Petroleum Maatschappij (Exploration and Production).
15. SINDOC(81)14). Risø National Laboratory, Roskilde, Denmark.
16. Senders, J. W. & Moray, N. P. (1991). *Human error. Cause, prediction, and reduction*. Hillsdale, N.J.: Lawrence Erlbaum Associates.
17. Shappell & Wiegmann. (2000). *The Human Factors Analysis and Classification System—HFACS*. FAA. US Department of Transportation, p. 2.
18. Shorrock, S., Young, M., Faulkner, J. (2005). Who moved my (Swiss) cheese? *Aircraft & Aerospace*, January/February 2005, 31-33.
19. Turner, B. A. (1978). *Man-made disasters*. London: Wykeham.

20. Williams, J. (1988). A data-based method for assessing and reducing human error to improve operational performance. In Proceedings of IEE Fourth Conference on Human Factors in Power Plants, Monterey, CA, 6-9 June.
21. Williams, J. (1996). Assessing the likelihood of violation behaviour: a preliminary investigation. Manchester: Department of Psychology, University of Manchester.
22. Yerkes, R. M. & Dodson, J. D. (1908). The relation of strength of stimulus to rapidity of habit-formation. *Journal of Comparative and Neurological Psychology*, 18, 459–482.

Page intentionally left blank

1. PREAMBLE

In the evening of the 1st July 2002, a Bashkirian Airlines TU-154 collided at flight level 350 with a DHL Boeing 757 cargo aircraft above the Lake Constance, Germany, near the city of Überlingen, resulting in 71 fatalities. Both aircraft were under the control of the Zurich ACC. Several attempts at the identification of the underlying human and organisational factors that led to this mid-air collision followed the accident, in particular the report by the German agency for air accident investigations (BFU).

On the 7th-8th September 2004, a two-day workshop was convened at the Eurocontrol Experimental Centre (EEC) in Brétigny (France) to review the final BFU report once more, and to consider whether new safety recommendations were appropriate for assuring the continued safety of European air traffic. The first day reviewed the normal defences that were defeated in this accident, with reference to the HFACS framework (see Annex 1). However, the workshop, particularly led by Professors James Reason and Erik Hollnagel, triggered a discussion about the relevance of the “Swiss Cheese Model” to analyze accidents such as the Überlingen disaster. Indeed, although it has been instrumental towards a systemic and organisational perspective on safety, the “Reason Model” is subject, as any model, to some limitations. In his contribution to the Brétigny workshop, J. Reason himself questioned the use of the model under this rather provocative title: “*Überlingen: Is Swiss cheese past its sell-by date?*”.

Following the Brétigny workshop, the EEC wished to further elaborate on this methodological aspect, and tasked a team of three¹ to review the issue. This report is the outcome of their work.

¹ James Reason, Erik Hollnagel, Jean Paries

2. SCOPE AND PURPOSE OF THE STUDY

It is now broadly recognized that accidents in complex systems occur through the concatenation of multiple factors, where each may be necessary but where they are only jointly sufficient to produce the accident. All complex systems contain such potentially multi-causal conditions, but only rarely do they arise thereby creating a possible trajectory for an accident. Often these vulnerabilities are “latent”, i.e. present in the organization long before a specific incident is triggered. Furthermore, most of them are a product of the organization itself, as a result of its design (e.g. staffing, training policy, communication patterns, hierarchical relationship,) or as a result of managerial decisions.

In 1990, James Reason (1990), then a professor with the University of Manchester, provided an crucial contribution to the concretization of this idea by proposing a “model” of how accidents could be seen as the result of interrelations between real time “unsafe acts” by front line operators and latent conditions. This model turned out to be highly pedagogical, and a large number of safety analysts around the world quickly started to use it in different industries. The ICAO Human Factors and Flight Safety Working Group adopted it in the early 90’S as a conceptual framework. As usual, many people – including the author himself – tried their own variants and refinements of the initial model. For example, in 2000 Shappell & Wiegmann adapted the Reason model to develop the Human Factors Analysis & Classification System (HFACS), an incident/accident analysis methodology sponsored by the Office of Aviation Medicine of the US Federal Aviation Administration.

While much of the accident investigation community swiftly adopted the Swiss Cheese Model (SCM), not least in the aviation domain, the enthusiastic use sometimes relied on interpretations of the model’s semantics that went rather far beyond what was initially intended. The aim of this report is therefore to discuss the relevance and limitations of using the SCM, particularly from an ATM accident investigation perspective. In his 1997 book *Managing the risks of organisational accidents*, Reason did warn that “the pendulum may have swung too far in our present attempts to track down possible errors and accident contributions that are widely separated in both time and place from the events themselves”. Yet despite this caution the use of the model continued to grow, so much that Shorrock & al (2005) remarked that “(i)ronically, it seems that the only person to question the use of Reason’s Swiss Cheese model is Reason himself!

Shorrock & al see this as an invitation to “pass a critical eye over the interpretation of Reason’s organisational approach and its application to accident investigation”. They question whether the focus on organisational psychopathology has been overplayed, and whether “we should redress some of our efforts back to the human at the sharp end”. Taking the example of en-route Airprox incidents analyses in the UK, they notice that in ATC, “very few organisational factors were mentioned in the published incident reports” and that “active errors and human performance at the sharp end [seem to be] currently the major factors contributing to incidents”.

Whether this is a deep feature of ATC incidents aetiology, or a surface by-product of the ATC incident analysts’ mindset, is obviously disputable. Ironically, the perspective of the Brétigny workshop was, on the contrary, that the Überlingen accident had no clear single causes or front line operator errors in the traditional sense. It was rather a confluence of a number of conditions, many of which were clearly sub-optimal, that led to this tragic event.

The core subject of the document is the Swiss Cheese Model (SCM). The purpose is to provide the reader with some means to comprehend fully the scope of the SCM, as well as to assess the appropriateness – or otherwise – of the criticisms that have been put forward. The purpose is also to discuss how the SMC has been used by practitioners, in order to assist these practitioners in understanding the SCM capabilities and limitations and hopefully prevent overly dogmatic implementations.

This document does not intend to discuss the specificity of ATC incidents concerning the role of front line operators “errors”, nor the fact that different domains may have a different sensitivity to, or a different resilience towards, sharp end errors, failures and the like.

The document includes the following next sections:

- Section 3 will present the history of the model, as experienced by J. Reason himself, hence provide an account of when, why and how the SCM was created and developed.
- Section 4 will discuss the conceptual status of the model, i.e. the various ways in which it can be interpreted and used.
- Section 5 will briefly review the dissemination of the model.
- Section 6 will review the published criticisms of the model.
- Section 7 will further elaborate on the conceptual status of the SCM, i.e., the various ways in which it can be interpreted and used.
- Section 8 will offer some conclusions.

3. HISTORY OF THE MODEL

The model had its origin in 1987-8 during the writing of *Human Error* (Reason, 1990). The original intention for the book was to provide an essentially cognitive psychological account of the nature, varieties, and the mental sources of human error. The underlying question was: What can the appearance of relatively non-random error forms tell us about the largely hidden processes that govern our thoughts and actions? The book was aimed primarily at J. Reason's peers: academic cognitive psychologists. There was no plan at the outset to include the chapter on latent errors and systems disasters in which the Mark 1 version of the model appeared. Its inclusion was prompted by two things: first, the abandonment of a lengthy chapter on the history of error studies, and the second (and far more important) influence was the spate of disasters occurring in the late 1970s and 1980s. These included Flixborough, *Challenger*, Three Mile Island, Bhopal, Chernobyl, the *Herald of Free Enterprise* and the King's Cross Underground fire.

Each of these accidents had been extensively investigated and the reports had made it clear that the performance of those at the sharp end (who may or may not have made errors, but mostly did) was shaped by local workplace conditions and upstream organisational factors. It became obvious that one could not give an adequate account of human error without considering these contextual system issues. There was nothing new in this observation—Barry Turner (1978) and Charles Perrow (1984), among others, had already given eloquent and persuasive descriptions of these systemic influences in the aetiology of major accidents.

Chapter 7 began by making a distinction between *active errors* and *latent errors*. The latter were likened to 'resident pathogens' in the body. Their adverse consequences could lie dormant within the system for a long time, only becoming evident when they combine with other factors to breach the system's defences. Such a notion was implicit in the Rasmussen and Pedersen (1984) discussion of the 'anatomy' of an accident, and Perrow had stressed the role of defences-in-depth in creating complexity, tight-coupling and opacity in the system. As J. Reason began writing the chapter he had in his head these two notions: the biological or medical metaphor of pathogens, and the central role played by defences, barriers, controls and safeguards (analogous to the body's auto-immune system).

All the major accidents listed above had occurred in complex productive systems. If we are to understand how humans contribute to the breakdown of these systems, we need first to identify the necessary and 'healthy' elements of production in order to describe how and why they might fail. John Wreathall and J. Reason depicted these as a sequence of five 'planes' lying one behind the other: top level decision makers, line management, preconditions, productive activities and defences. Failures can arise at anyone of these levels. The benign and the pathological aspects of each plane are shown in Figure 1 and Figure 2. Note that there is nothing Swiss-cheese-like in anything but the last plane where there is a single hole. J. Reason did not originate the 'Swiss cheese' label—that was probably Rob Lee, then Director of the Bureau of Air Safety Investigation (BASI) in Canberra.

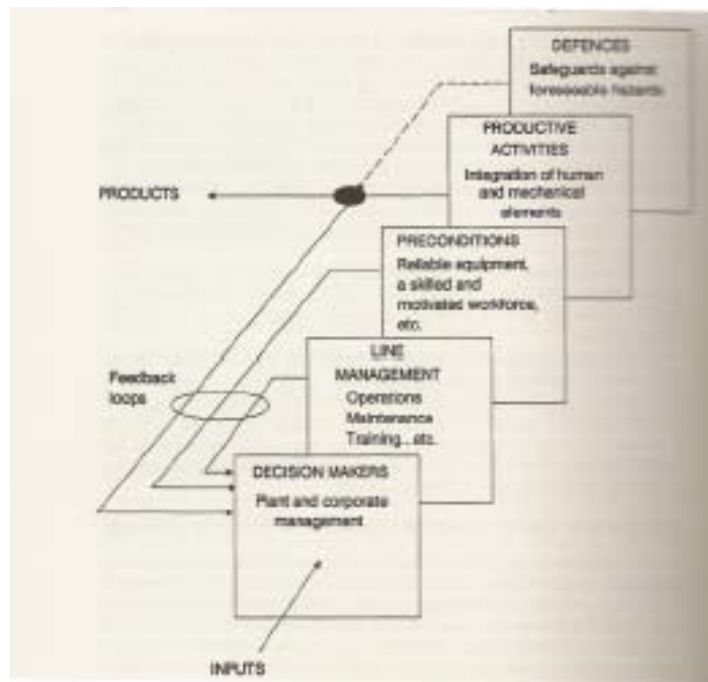


Figure 1

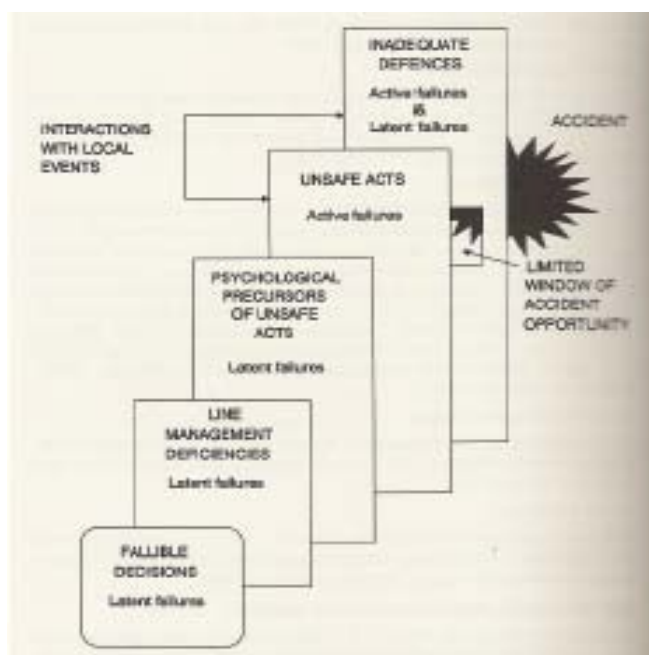


Figure 2

However, a subsequent representation in Chapter 7 (Figure 3) takes on a distinctly 'cheesy' flavour. It shows an accident trajectory passing through successive 'slices'. Its purpose was to show the dynamics of accident causation arising from interactions between latent failures and a variety of local triggering events.

One can suspect that it is this picture which gave rise to the Swiss cheese label, even though most applications of the model (e.g., HFACS, ICAM, etc.) derive from Figure 2, though with ‘holes’ added. It is interesting to think that had the author(s) represented these elements of production and their corresponding weaknesses as boxes rather than ‘planes’ (as in the Mark II model) there would probably have been no ‘Swiss cheese’ metaphor, or at least not almost immediately as actually happened.




Figure 3

The Mark II model (see Figure 4) was developed in the early to mid-1990s. It reduced the four productive planes (in Figure 1 and Figure 2) to three (organisation, workplace, person), but extended the single defensive layer to three layers. The aim here was to allow more specificity with regard to the influences at each level. J. Reason also distinguished errors and violations and their corresponding provocative factors—this was made possible by the work of Jerry Williams who, in HEART (1988), had already identified the relative impact of various error-producing factors; but then spent a sabbatical year in Manchester doing the same for violation-producing factors (Williams, 1996). The organisation box now included corporate culture and organisational processes as well as management decisions. This was a time when the impact of safety culture was becoming more and more evident, and J. Reason’s work with Shell and other organisations on proactive process measurement had shown how generic system activities (designing, building, operating, maintaining, managing, budgeting, scheduling and the like) determine an organisation’s ‘safety health’ or resilience. The final addition was a separate latent failure path leading from the organisation box to the defences. This pathway accommodated the fact that in many disasters—King’s Cross, *Challenger*, *Piper Alpha*—there were no proximate active failures, simply long-standing systemic pathogens.

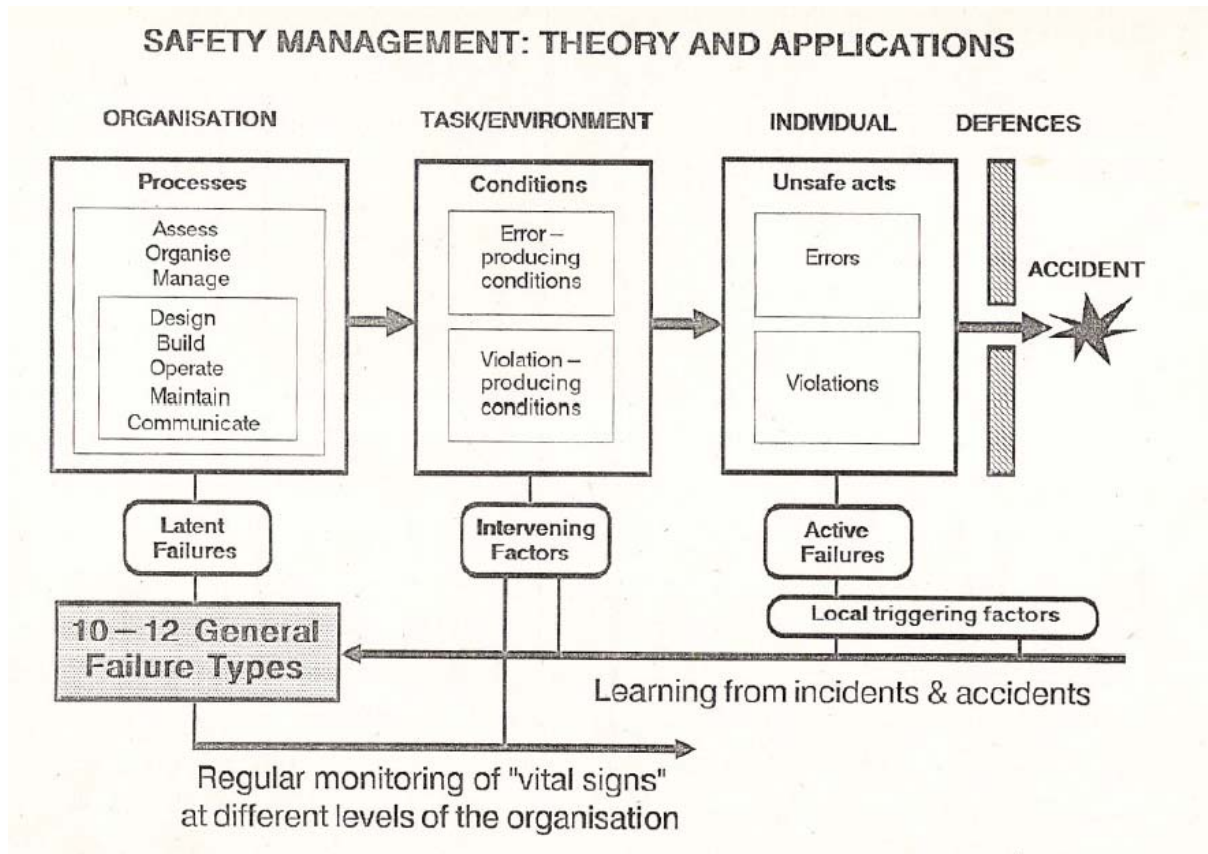


Figure 4

The Mark III version of the model (see Figure 5) appeared in *Managing the Risks of Organizational Accidents* (Reason, 1997). There were a number of significant changes.

- It started with the premise that any model of accident causation must have three basic elements: hazards, defences and losses.
- The 'planes' are now represented as undisguised Swiss cheese slices, but they are not labelled. They include all the many barriers, defences, safeguards and controls that any given system might possess. This was felt necessary because it is often very hard to distinguish productive and protective system elements. Many sub-systems serve both goals (e.g. pilots, ATCOs, train drivers, etc.).
- An important addition was an explanation of how the holes, gaps or weaknesses arise. Short-term breaches may be created by the errors and violations of front-line operators. Longer-lasting and more dangerous gaps are created by the decisions of designers, builders, procedure writers, top-level managers and maintainers. These are now called *latent conditions* rather than latent errors or latent failures. A condition is not a cause, but it is necessary for a causal factor to have an impact. Oxygen is a necessary condition for fire; but its cause is a source of ignition. The use of this term allows us to acknowledge that all top-level decisions seed pathogens into the system, and they need not be mistaken. Allocating resources between departments is rarely done by giving out equal shares; some departments get more than others for are judged to be sensible reasons at the time. But those with smaller slices of the resource cake will often have poorer equipment, extra time pressure, under manning and other error-provoking factors. The existence of latent conditions is a universal in all organisations, regardless of their accident record.

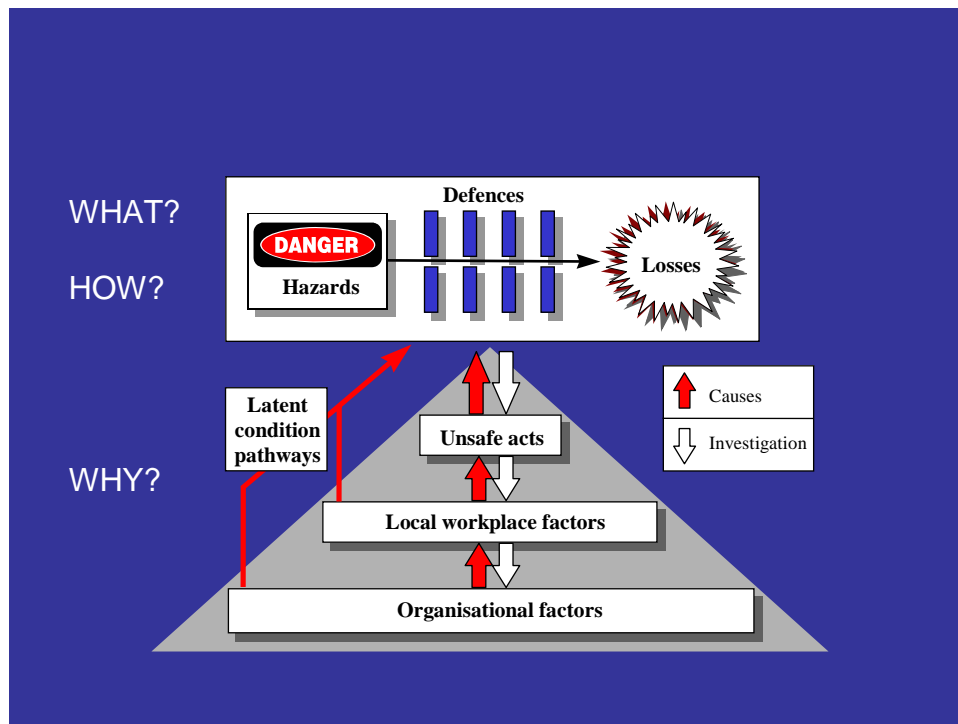


Figure 5: the Mark 3 version

4. FUNCTIONS OF THE MODEL

One possible reason for the different views on the models merits is that it can be used to serve three different purposes. These are discussed in the following.

4.1. SCM AS A CONCEPTUAL FRAMEWORK

The SCM is a heuristic explanatory device for communicating the interactions and concatenations that occur when a complex well-defended system suffers a catastrophic breakdown. In particular, it conveys the fact that no one failure, human or technical, is sufficient to cause an accident. Rather, it involves the unlikely and often unforeseeable conjunction of several contributing factors arising from different levels of the system. It also indicates what defines an *organizational accident*, namely the concurrent failure of several defences, facilitated, and in some way prepared, by sub-optimal features of the organisation design.

In this regard it has proved very successful. It is a simple metaphor—easily remembered and passed on—that encompasses what is often a very complex story. A Google search on ‘Swiss cheese model of accidents’ yielded around 18,400 hits covering a wide range of hazardous domains. Many of these involve passing on the model to various professional communities. A high proportion of these messages are aimed at health carers. We will discuss the dissemination of the model shortly.

4.2. SCM AS A MEANS OF COMMUNICATION

The SCM also acts as a framework for accident investigation. HFACS (Shappell & Wiegmann, 2001) is one of the most widely used of these derivative techniques. Other examples include BHP’s ICAM, Shell’s Tripod Beta, and various so-called root cause analysis techniques—though there is no notion of ‘root cause’ in the Mark III Swiss cheese model. J. Reason’s definition of a ‘root cause’ is the contributing factor that you are working on when the money or the time runs out.

HFACS is derived largely from the Mark I model (as depicted in Figure 1 and Figure 2), though the layers have been somewhat modified (see Figure 6). Shappell and Wiegmann’s most important contribution is the degree to which they operationalised the application of the model so that it can be used by a wide range of investigators. They criticise the original model for failing to identify the cheese holes more precisely. But such specificity was never the original intention. The model was intended to be a generic tool that could be used in any well-defended domain—it is for the local investigators to supply the local details.

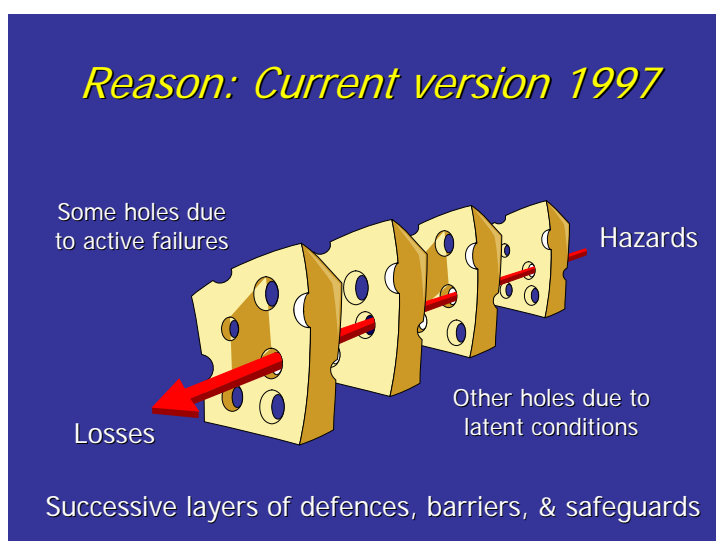


Figure 6

4.3. SCM AS A BASIS FOR ANALYSIS

The model has also been applied to proactive process measurement—the repeated assessment of a limited set of ‘vital signs’ that collectively give some indication of the current state of ‘safety health’ and the factors that are most in need of remediation. The first of these tools was Tripod-Delta created for Shell in 1988-1990 by Wagenaar, Hudson, Reason, Benson and Groeneweg. An early depiction of the philosophy underlying this technique is shown in Figure 7. This philosophy is discussed in detail in Reason 1990a, 1990b, 1997). Other prospective techniques based upon these ideas are REVIEW (created for British Rail), MESH (tailored for British Airways Engineering and for Singapore Airlines Engineering Company) and PAOWF (a technique for measuring ‘leading indicators’ in the US nuclear power industry).




Figure 7

5. DISSEMINATION OF THE MODEL

It is said that Freud discovered that he was famous when he found his cabin steward reading the *Psychopathology of Everyday Life* on his first trip to the US in the early 1900s. In the late 1990s, J Reason was taken around the air traffic control tower by Vancouver Harbour where he was introduced to a young ATCO. When he heard his name he said, 'ah yes, you're the Swiss cheese man'. There is no doubt that the Swiss cheese model spread very widely very fast. Some of the agents of this dissemination are listed below.

Human Error has been translated into five languages and its annual sales continue to rise. *Managing the Risks of Organizational Accidents* (a much more readable book aimed at practitioners) has also been translated into Japanese, and it continues to sell. Since 1990, J.Reason has given more than 280 presentations relating to the model in Europe, the US, Canada, Mexico, the Middle East, Africa (Nigeria, Gabon, Ethiopia), South-east Asia, Japan, Hong Kong, Australia and New Zealand (most on many repeated visits). The demand for next year is as strong as ever. Much more significant, however, has been the work of some very influential agents. Some of these are listed below.

- Dr Rob Lee, Director of BASI insisted that his investigators use the 'Reason model'.
- Captain Dan Maurino, ICAO Human Factors Digest No.7 (1993) introduced the Mark I model to the aviation world. He and Rob Lee were also influential in getting a form of the model adopted by ICAO Annexe 13 (8th edition).
- Jean Paries was also influential in introducing the Mark I model to the aviation world, particularly for accident investigations.
- Professor Jan Davies, Department of Anaesthesia, University of Calgary. She and her colleagues applied the model to an anaesthetic fatality in 1991.
- Roger Taylor, British Rail, now Rail Safety & Standards Board.
- The US Nuclear Regulatory Commission and the Institute for Nuclear Power Operations (Atlanta).
- Various NASA managers, scientists and astronauts.
- Peter Harle of the Transport Safety Board (Canada).
- CASA (Australia).
- Captain Jeremy Butler, British Airways, now a non-executive board member of the National Patient Safety Agency.
- Captain Bertrand de Courville, Air France.
- Dr Lucian Leape, patient safety pioneer at Harvard School of Public Health.
- Dr Don Berwick, President, Institute of Health Improvement, Boston.
- Depart of Health (UK): *An Organisation with a Memory*.
- US Institute of Medicine: *To Err is Human*.

6. PUBLISHED CRITICISMS OF THE MODEL

The first main external criticism has been that the model is insufficiently specific regarding the nature of the holes in the cheese and their inter-relationships. Thus, it is not easily applicable as an investigation tool. Thus Luxhoj & Kauffeld (2003) writes that:

One of the disadvantages of the Reason model is that it does not account for the detailed interrelationships among causal factors. Without these distinct linkages, the results are too vague to be of significant practical use.

Dekker (2002, p. 119-120) adds that:

The layers of defence are not static or constant, and not independent of each other either. They can interact, support or erode one another. The Swiss cheese analogy is useful to think about the complexity of failure, and, conversely, about the effort it takes to make and keep a system safe. It can also help structure your search for distal contributors to the mishap. But the analogy itself does not explain:

- *where the holes are or what they consist of,*
- *why the holes are there in the first place,*
- *why the holes change over time, both in size and location,*
- *how the holes get to line up to produce an accident.*

This is up to you, as investigator, to find out for your situation.

Finally, Shappell & Wiegmann (2000) notes that:

In many ways, Reason's 'Swiss cheese' model of accident causation has revolutionized common views of accident causation. Unfortunately, however, it is simply a theory with a few details on how to apply it in a real-world setting. In other words, the theory never defines what the 'holes in the cheese' really are, at least within the context of everyday operations.

A second main external criticism echoes J. Reason's own feelings that "the pendulum may have swung too far in our present attempts to track down possible errors and accident contributions that are widely separated in both time and place from the events themselves". It is best worded by Steve Shorrock & al.

Reason (1990) stated that "systems accidents have their primary origins in the fallible decisions made by designers at high-level (corporate or plant) managerial decisions makers". Active errors were therefore seen as symptoms or token of a defective system. It apparently became the duty of accident investigators and researchers to examine the psychopathology of organisations in the search of cues.

[...]

One implication of the organisational approach has been the tenacious search for latent conditions leading up to an accident. There are serious flaws in such prescriptive implementation. Whilst the importance of analysing human factors throughout the accident sequence is not in question, the dogmatic insistence on identifying the latent conditions could and should be challenged in case where active failures played a major part.

[...]

We address the question of whether the focus on latent errors has become too strong, and whether we should redress some of our efforts back to the human at the sharp end.

[...]

The mapping between organisational factors and errors or outcomes, if any such mapping can be demonstrated with an appropriate degree of certainty, is complex and loosely coupled. However, when analysing events and conditions in an accident evolution, the SCM makes it tempting to draw a line from an outcome to a set of 'latent conditions'. This invites 'hindsight bias' [...].

Actually, Shorrock & al raise several issues:

- active errors may be the dominant factor: latent conditions are clearly important, but sometimes people really just slip up,
- the causal link, or even the connection, between distant latent conditions and accidents are often tenuous, and only visible with the benefit of hindsight,
- latent conditions can always be identified, with or without an accident,
- some latent conditions may be very difficult to control, or take many years to address,
- misapplication of the model can shift the blame backwards, from a 'blame the pilot' to 'blame the management' culture,
- highlighting management problems may hide very real human factors issues, like the impact of emotion on performance, and hamper the research needed to better understand human fallibility.

A few comments can be made on these critiques. The risk of shifting the blame towards the managers has been clearly acknowledged by J. Reason and newer versions of the model do not refer to 'unsafe decisions' or managerial failures but rather to organisational features. The fact that front line operators' slips sometimes fully accounts for the accident scenario does not mean that it explains the accident from a safety management perspective, and that 'fixing' the operator therefore is the right safety management strategy. The fact that deterministic causal connection between latent conditions and accidents cannot easily be identified (particularly before the event), does not rule out that efficient prevention policy can be based on addressing latent conditions. Although such connections may be long and difficult to control, they may also offer a real opportunity for effective accident prevention. Hindsight is a problem if it is used as a basis for holding individuals responsibilities (which seems to remain a subliminal preoccupation for Shorrock & al). From a safety management perspective, the key point is to identify, as well as possible, the potential contributors to a multi-factorial process. Here hindsight can be of benefit, although it should be used with care.

From an academic perspective the weakness of the SCM is that it is irrefutable by standard scientific methods – a distinction it shares with many other models from Freud's to Endsley's. There is no 'crucial experiment' that can be used to reject the model. But then it was never intended to be a scientific model on that level. It may therefore be unfair to level criticisms such as:

And although the emphasis is upon the promotion of active failures by latent preconditions, the model lacks the fully organizational level of organizational analysis needed to describe inter-organizational phenomena, or to fully close the critical gap between latent and active accident precursors in theoretical terms.

(Pidgeon, 1977)

Or:

The Reason swiss cheese is no longer adapted to anticipate today's accidents.

Nancy Leveson (2004)

7. FURTHER COMMENTS ON THE SWISS CHEESE MODEL (SCM)

The preceding sections have provided an account of how the Swiss Cheese Model (SCM) was developed and also made clear that it can be interpreted in several ways. Some of the published criticisms of the SCM have also been presented. The purpose of this section is to comment on what may be called the conceptual status of the SCM, i.e., the various ways in which it can be interpreted and used. This is necessary to comprehend fully the scope of the SCM, as well as to assess the appropriateness – or otherwise – of the criticisms that have been put forward and the way in which it has been used by practitioners.

7.1. THE STATUS OF THE SCM AS A MODEL

When referring to models, in accident descriptions as in behavioural sciences in general, the term is often used rather loosely. A model is essentially a simplified representation of something else, of phenomenon or event such as an accident or of a system such as an organisation or ATM. The simplification makes it easier to understand the essential features of that which is modelled. A model can be used for a number of different purposes, such as a means of communication (for which an analogy or metaphor may also serve the purpose), as a replacement for the real-world phenomenon or system (where the model is easier to manipulate and study, e.g., as a scale model of an airplane in a wind tunnel), or as an instantiation of some basic principles or “laws” (for instance a formal decision model such as Elimination-by-Aspects or a model of psychomotor behaviour such as Fitts’ Law).

Two widely different uses of models are the *retrospective*, where the model is the basis for explaining or understanding something, and the *prospective*, where the model is the basis for predicting something – including measurements of present states as an indicator of possible future states. A model is thus a convenient way of referring to the shared sets of axioms, assumptions, beliefs, and facts about a phenomenon that makes it possible to form an understanding of what goes on and to make predictions about what will happen. In science, the predictive capabilities are often considered to be of primary importance, since the predictions allow the underlying hypotheses, hence the model, to be falsified.

Although the “Swiss cheese” normally is referred to as a ‘model’, that appellation is only appropriate for some of the versions. For practical reasons we will nevertheless continue to refer to it as the SC Model. It is in reality less important what it is called, than what it can be used for.

As pointed out earlier in this report, the SCM has been used for three different purposes. The first is as a heuristic explanatory device (communication); the second is as a framework for accident investigation (analysis); and the third is as a basis for measurements (measurements). Each of the three purposes will be considered in the following.

7.2. THE SCM AS A MEANS OF COMMUNICATION

The first version of the SCM (Figure 8) provides a way to think of how accidents happen (a heuristic explanatory device), or in other words a specific accident model (cf. Hollnagel, 2004). The “logic” is that accidents are the result of a sequence of events or a serial development. This is similar to the principle of Heinrich’s Domino model (Heinrich, 1931), although it is expressed in a less rigid and absolutistic manner. Another important feature of the SCM is that the events in the sequence are described as *failures* at different organisational levels, going from senior management to unsafe acts at the sharp end. This version of the SCM can therefore also be seen as representing the kind of thinking that later became known as the *sharp end, blunt end* descriptions of accidents.




Figure 8: SCM Mark I (pathological aspects)

7.3. NORMATIVE MODEL OF ORGANISATIONS

The conceptual ancestry of this version of the SCM is a normative model of what an organisation is. Organisations, whether they are a result of natural evolution or design, generally function in a hierarchical fashion. This means that actions, decisions and directives made at a higher level are passed on to a lower level, where they either are implemented directly, or interpreted in some way before they are passed on to the next level below, etc. The basic principle of organisational control is simply that higher levels control what happens at lower levels, although control more often is in terms of goals (objectives) and criteria than instructions that must be carried out to the letter.

The normative model of organisational functions is easily seen if instead of Figure 8 we use the benign version of the Mark I SCM (Figure 1). The result is then: Plant and corporate management decisions, line management, preconditions (for work), productive activities, and defences. It is, indeed, this normative model that is the very basis for the principle used to explain accidents as failures at anyone of these stages: management decisions propagate downwards and progressively turn into productive activity; Bad management decisions propagate downwards and progressively turn into unsafe activity, and possibly accidents.

Real life may, however, often differ from the norm and this is not least the case for organisations. On way in which this is recognised is the distinction between the formal and informal organisation. The normative model of an organisation can therefore not be taken as universally valid. The actual way in which something happens may differ from what is prescribed – and normally does – although the outcome may still be the intended one. Similarly, the “historical” account of a given event may also be more complex than the formal model, for instance because decisions at the higher levels may come about as a response to events at the sharp end, such as a revision of procedures after an accident.

7.4. ORDERLINESS AS AN ARTEFACT OF RETROSPECTION

According to the normative model of an organisation events necessarily take place in a sequence or in serial order. Indeed, this is the very principle of hierarchical control. There are, however, two problems with this description. One, that events in practice may happen in a different order and that this order may be hard to predict. Second, that the *normative serialisation* may be confused with the *temporal orderliness* of events that have happened.

It is a basic fact that any description of events in the past (e.g., an accident), will show them as ordered in time: one event A will necessarily have happened before another event B. Although we in some cases may be satisfied by saying that two events happened at the same time, accidents analyses typically make an effort to establish a clear temporal order of events as they took place, particularly in the last hours, minutes, or seconds before the accident (i.e., before the final event that marks the accident). This orderliness is, however, an artefact of retrospection.

We have known, at least since the days of David Hume, that causes must be prior to effects, e.g., that A must happen before B. But we also know that the temporal orderliness of two events does not mean that A necessarily is the cause of B. Such a conclusion is logically invalid and furthermore disregards the role of coincidences.

The correspondence between normative serialisation and temporal orderliness in the SCM, and in the graphical rendering, is an unfortunate but entirely unintended consequence. That one event (a line management decision) happens prior to another event (an action at the sharp end), is in one case the consequence of organisational logic and in the other of the temporal ordering of events. As an accident model, the SCM therefore illustrates how an accident *could* happen, but not how it *must* happen. The latter is only the case if organisations function in strict accordance with the rules and principles laid down, which in practice never happens.

In summary, the strong point of the SCM as a communication tool is that it has been instrumental in developing the understanding of accidents as the outcome of failures at several stages, as a (complex) combination of active failures and latent conditions, rather than as the result of isolated events at the sharp end. It is thus a powerful vehicle for explanations. The weak point is that this rendering may be interpreted to mean that accidents result from a specific sequence of failures, although that is neither the case nor the intention. As the Mark I and Mark II versions of the SCM only represent a single type of accidents, they are not as such form appropriate for making predictions.

7.5. PROXIMAL AND DISTAL CAUSES

As an accident model, the SCM reinforced the already existing view that the proximate causes of accidents were to be found in failures at the sharp end, specifically in “human errors” (cf. Rasmussen, et al., 1981; Senders & Moray, 1991). This view is part and parcel of many commonly used methods such as HFACS or HERA. Although the SCM emphasises the need to understand the complexity of the situation and in particular the role played by latent conditions, the accident analysis nevertheless reverses the path through the normative organisation model described above. An alternative is represented by a view of accidents as due to a combination of human, technological and organisational factors. (In the Nordic countries this is known as the MTO-model, from the triad *Människa* (Man) – *Teknik* (Technology) - *Organisation*.) The MTO thinking was introduced by the Swedish Nuclear Inspectorate in the early 1980s, but has only in recent years become internationally recognised.

Although the SCM does include both technological and organisational factors, hence nominally corresponds to an MTO-model, the difference lies in the order in which these factors is addressed. An MTO-based approach, such as CREAM (Hollnagel, 1998), considers the three groups of equal importance and can therefore begin the analysis with any of them. In contrast, methods such as HFACS and HERA begin by looking at human failures and considers technological and workplace factors as secondary influences. (This, by the way, is consistent with the established practice in most HRA methods.)

7.6. BARRIERS AND DEFENCES

As an accident model, the SCM describes accidents as a combination of specific events and the failure of one or more barriers – or of *all* barriers if they are serial rather than parallel. The events are described in terms of the normative organisation model, to which is added a set of barriers. In the first renderings of the model (SCM Mark II), barriers were shown only at the sharp end. This is not surprising since an accident investigation for natural reasons starts with the proximate events, including barriers and defences that somehow failed to meet their objectives. Although it is not always expressed directly, the SCM logically supports the view that (failed) barriers can be found at any level of the organisation or – what is essentially the same thing – at any stage of the developments that led to the accident. This is consistent with the view that “everybody’s blunt end is somebody else’s sharp end”. The positioning of the barriers at the sharp end is therefore a symbolic simplification and should not be taken literally.

Newer versions of the SCM (cf. Figure 6) have been modified to emphasise the role of (failed) barriers or defences and to deemphasise the normative organisation model. An additional consequence is that the emphasis on human failures at the sharp end also is reduced. In this case the existence of hazards is taken for granted without trying to account in detail for how they may have come about. The SCM instead puts the focus on the various barriers that may exist in the system, and how deficient barriers may fail to prevent a hazard from resulting in a loss.

As the basis for accident analysis, the SCM is valuable but incomplete. It is valuable because it makes clear that accidents have complex causes (or explanations) and because it brings forward the effects of factors that may otherwise be hidden from view (i.e., latent conditions). The Mark III version of the SCM furthermore relaxes the requirement to the types of barriers in a system and to their specific order. Although a relaxation of the SCM structure broadens its applicability, it also offers less support for a structured analysis. The SCM also focuses on barriers rather than hazards. This may be taken as representing the view that it is more efficient to prevent accidents by strengthening system barriers than by eliminating causes. From a contemporary perspective it is, indeed, a considerable strength of the SCM that it avoids the notion of root causes, despite the fact that many practitioners seem to relish this concept. A final comment is that the earlier versions of the SCM gave higher priority to human action failures than to organisational or technical failures. This has led to a bias in some of the methods associated with the SCM Mark I and Mark II, in contrast to the more impartial perspective offered by a contemporary systemic view of accidents.

7.7. CONTEMPORARY VIEW OF ACCIDENTS

The understanding of how accidents occur has during the last eighty years or so undergone a rather dramatic development. The initial view of accidents as the natural culmination of a series of events or circumstances, which invariably occur in a fixed and logical order (Heinrich, 1931), has in stages been replaced by a systemic view according to which accidents result from an alignment of conditions and occurrences each of which is necessary, but none alone sufficient (e.g., Bogner, 2002).

Indeed, it may even be argued that the adaptability and flexibility of human performance is the reason both for its efficiency and for the failures that occur, although it is rarely the cause of the failures. In that sense even serious accidents may sometimes happen even though nothing failed as such.

Adopting this view clearly defeats conventional accident models, according to which accidents are due to certain (plausible) combinations of *failures*. This is the logic of functions as represented, e.g., by the fault tree. But the fault tree only shows representative accidents. The more unusual accidents cannot be captured by a fault tree, one reason being that there are too many conjunctive conditions. What we see in accidents is that confluences occur, and predictive accident models must therefore not only recognise *that* confluences occur but also provide a plausible explanation of *why* they happen. If we relax the requirement that every accident *must* involve the failure of one or more barriers, the inescapable conclusion is that we need accident analysis methods that look equally to individual as to organisational influences. In other words, models of “human error” and organisational failures must be complemented by something that could be called socio-technical or systemic accident models.

The Swedish MTO line of thinking mentioned above is a case in point. It promotes a view of accidents as due to a combination of human, technological and organisational factors, related to performance variability. Performance variability management accepts the fact that accidents cannot be explained in simplistic cause-effect terms, but that instead they represent the outcome of complex interactions and coincidences which are due to the normal performance variability of the system, rather than actual failures of components or functions. (One may, of course, consider actual failures as an extreme form of performance variability, i.e., the tail end of a distribution.) To prevent accidents there is therefore a need to be able to describe the characteristic performance variability of a system, how such coincidences may build up, and how they can be detected. This reflects the practical lesson that simply finding one or more “root” causes in order to eliminate or encapsulate it is inadequate to prevent future accidents. Even in relatively simple systems new cases continue to appear, despite the best efforts to the contrary.

7.8. THE SCM AS A BASIS FOR ANALYSIS

The SCM has been used as the basis for analysis by promoting the view of accidents as a combination of specific events and a failure of barriers (cf. Figure 9). What the SCM more specifically shows is the causal – or temporal – ordering of deficiencies that can explain the accident. Relating to the above discussion, this need not be a necessary condition for the accident to happen in a general sense, but it is seen as a sufficient condition for the explanation of a specific accident, hence as a reference for accident analysis.

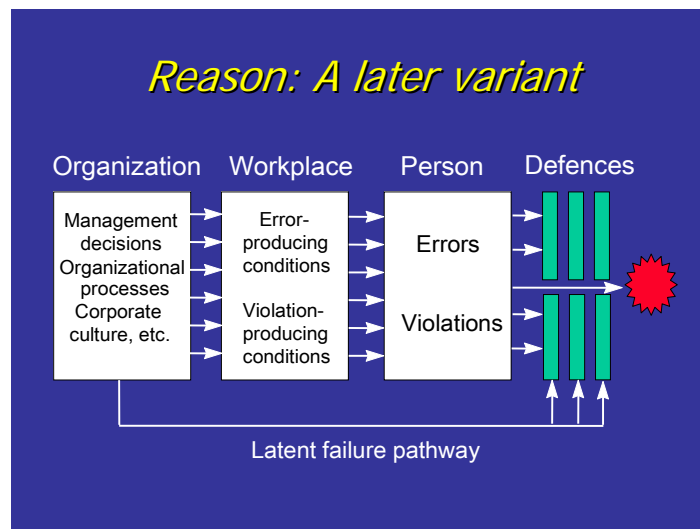


Figure 9: SCM Mark II

In this version of the SCM an accident is seen as the consequence of a series of deficiencies. This represents the thinking of the times, i.e., that accidents being negative events had to have causes that also were negative events. In other words: abnormal outcomes must be caused by abnormal behaviours, at different levels within the organisation. From the vantage point of the present, this view is no longer valid, since we now realise that accidents can arise from normal behaviours or normal events, or rather from the normal variability of established work practice. This is also noted by Pidgeon (1977):

Finally, its emphasis upon the notion of error and failures as the principal components in the pre-accident chain is problematic both in definitional as well as in philosophical terms, and forecloses to some extent the question raised in the first edition of Man-Made Disasters of the unintended consequences of otherwise ordered and 'rational' processes in ambiguous and ill-structured operating environments.

Pidgeon, 1997

7.9. SCM AS A BASIS FOR MEASUREMENTS

In addition to supporting explanation and analysis, the SCM also supports prediction in the sense of proposing a limited set of vital signs that provide an indication of the current state of “safety health”. This is consistent with the premise of the SCM, namely that accidents would not happen unless there were latent weaknesses in the system. These weaknesses may exist in the barriers and defences as well as in the established practices of work. The nearly irrefutable logic is that by identifying the potential weaknesses of a system in advance, it may be possible to intervene before an accident takes place.

Predictions based on models can differ with regard to precision and certainty. Some models are weakly predictive, in the sense that they generate predictions that are highly certain (i.e., likely to be true) but with limited precision. Other models are strongly predictive, in the sense that they generate predictions that are very precise but with limited certainty. An extreme example of the former is the stress-arousal model, also known as the inverted U-curve (Yerkes & Dodson, 1908), while a similar example of the latter is the HCR (Hannaman et al., 1984).

The SCM is best characterised as a weakly predictive model. It works by proposing a small number (10-12) of general failure types. These then serve as the basis for various measurements, which can be combined to express both the general “health” of a system and used to identify specific remedial actions. The prediction is thus of the general likelihood that an accident may happen, but not of where and when it may happen. Attempts to make predictions of the latter type have, on the whole, met with limited success.

The SCM is clearly less ambitious than other prediction methods, which often aim directly to identify likely human failures or even the probability by which these may occur. While such methods have been widely used the results are by many considered with a measured degree of scepticism, and by some with outright scorn. Most of these methods, particularly those that belong to the so-called first generation of HRA, are demonstrably based on rather shaky theories and models. In contrast to that, the SCM provides a robust model although with a low level of resolution. Rather than trying to predict the likelihood of specific types of accidents, it uses extensive practical experience to point to system functions that, if damaged, often are associated with the occurrence of incidents and accidents. This clearly has a considerable pragmatic value. The weak point is that the SCM does not propose a more detailed account of how latent failures, intervening factors, and active failures may combine, i.e., that it does not provide a more detailed accident model. It therefore does not fully support the contemporary view of accidents as the result of a conjunction or aggregation of conditions, none of which need to represent a failure as such.

8. CONCLUSIONS

Accidents come in many sizes, shapes and forms. It is therefore naïve to hope that one model or one type of explanation will be universally applicable. Some accidents are really simple, and therefore only need simple explanations and simple models. Some accidents are complex, and need comparable models and methods to be analysed and prevented. The developments in accident models, from the Domino model and onwards, has indeed been forced by cases that defied the current ways of thinking, almost in the nature of a Kuhnian paradigm shift. We are gradually faced with the need of models that can account for accidents as due to conjunctions of variability, both analytically and as predictions. Yet the introduction of a new model does not necessarily mean that those already existing become obsolete. It rather serves to highlight their strengths and weaknesses and thereby, in a sense, to ascertain when they should be used and when not.

From this perspective, and from the above comments, the SCM has an indisputable value as a means of communication, as a heuristic explanatory device. It has had a significant influence on the understanding of accidents, and therefore also on the practical approach to accident analysis and prevention. It has successfully been applied as a means of accident analysis and proactive measurements, although the level of resolution does not go as far as for other models. On balance this is, however, not necessarily a disadvantage since increased precision of analysis often is associated with reduced certainty of results.

The SCM does not provide a detailed accident model or a detailed theory of how the multitude of functions and entities in a complex socio-technical system interact and depend on each other. That does not detract from its value as a means of communication, but may limit its use in analysis and as a support for proactive measurements. On the other hand, the SCM was never developed to be a detailed model in this sense, and it is therefore hardly justified to criticise it for that. Indeed, as most of the criticisms fail to distinguish among the three primary uses of the SCM, they may be accusing it for failing to achieve something it never intended.

Page intentionally left blank

ANNEXE

Page intentionally left blank

ANNEX A - UBERLINGEN ACCIDENT ANALYSIS USING THE REASON MODEL FRAMEWORK

