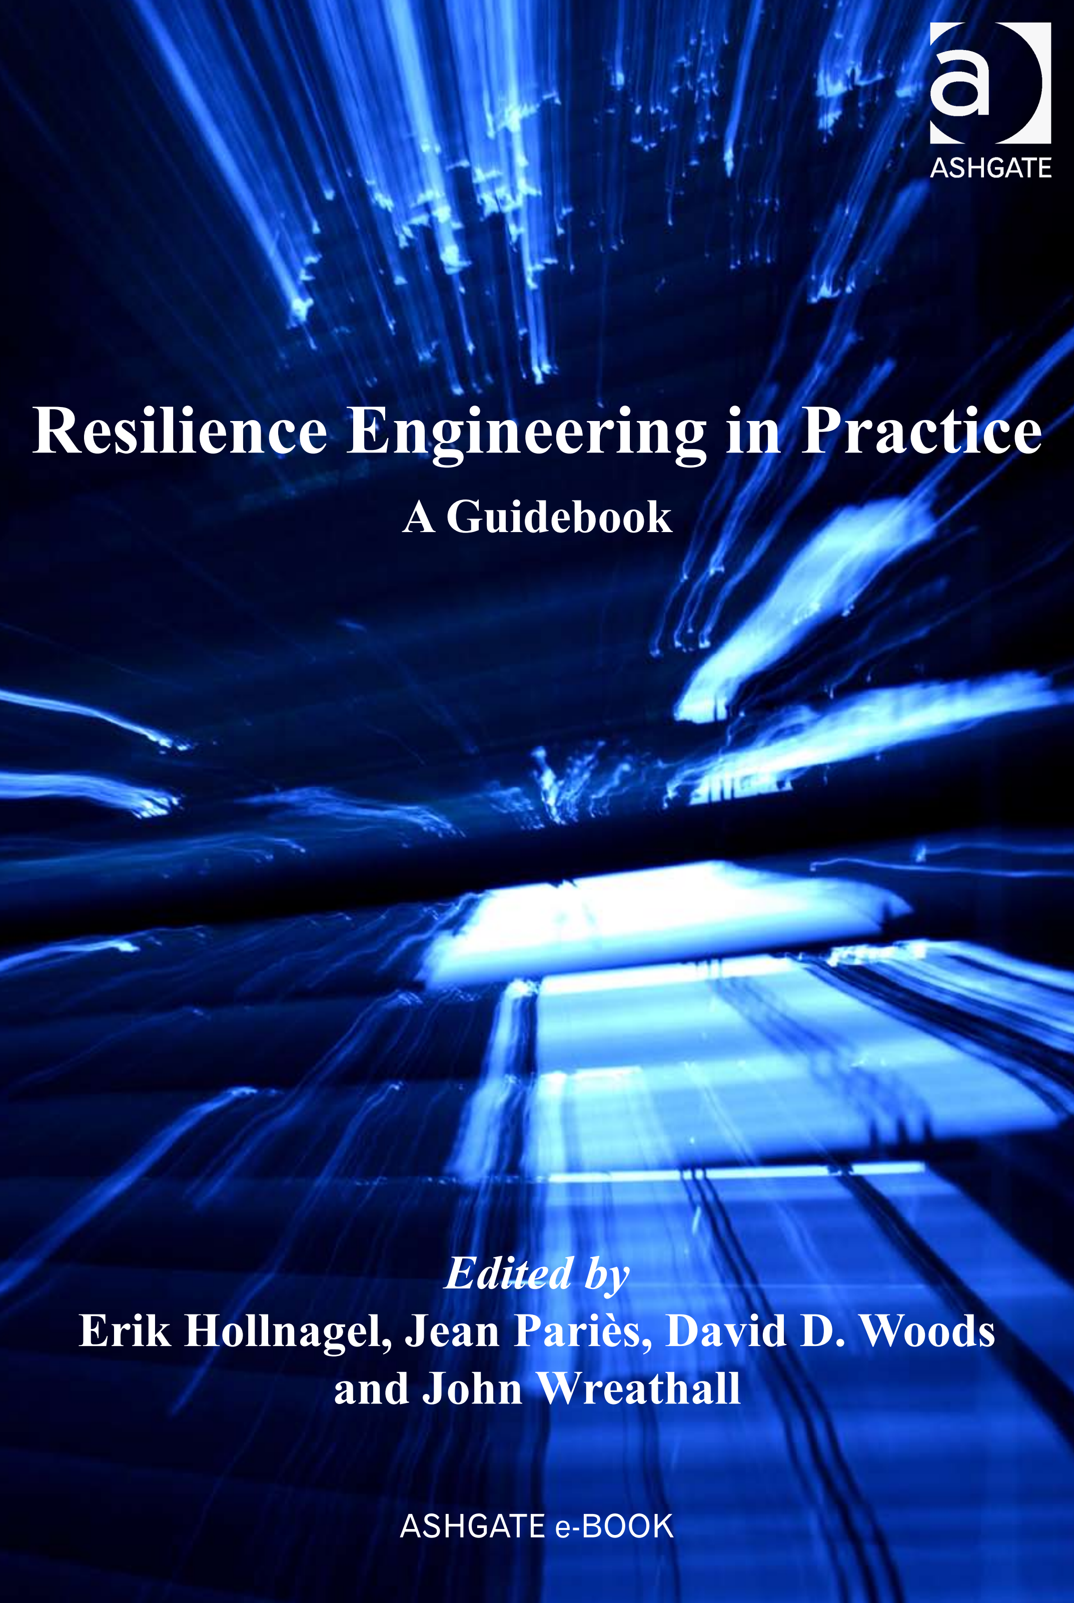

ASHGATE

Resilience Engineering in Practice

A Guidebook

Edited by

Erik Hollnagel, Jean Pariès, David D. Woods and John Wreathall

ASHGATE e-BOOK

RESILIENCE ENGINEERING IN PRACTICE

Ashgate Studies in Resilience Engineering

Resilience engineering has become a recognized alternative to traditional approaches to safety management. Whereas these have focused on risks and failures as the result of a degradation of normal performance, resilience engineering sees failures and successes as two sides of the same coin – as different outcomes of how people and organizations cope with a complex, underspecified and therefore partly unpredictable environment.

Normal performance requires people and organizations at all times to adjust their activities to meet the current conditions of the workplace, by trading-off efficiency and thoroughness and by making sacrificing decisions. But because information, resources and time are always finite such adjustments will be approximate and consequently performance is variable. Under normal conditions this is of little consequence, but every now and then – and sometimes with a disturbing regularity – the performance variability may combine in unexpected ways and give rise to unwanted outcomes.

The Ashgate Studies in Resilience Engineering series promulgates new methods, principles and experiences that can complement established safety management approaches. It provides invaluable insights and guidance for practitioners and researchers alike in all safety-critical domains. While the Studies pertain to all complex systems they are of particular interest to high-hazard sectors such as aviation, ground transportation, the military, energy production and distribution, and healthcare.

Resilience Engineering in Practice

A Guidebook

EDITED BY

ERIK HOLLNAGEL
MINES ParisTech, France

JEAN PARIÈS
Dédale SA, France

DAVID WOODS
Ohio State University, USA

JOHN WREATHALL
John Wreathall & Co., USA

ASHGATE

Published by
Ashgate Publishing Limited
Wey Court East
Union Road
Farnham
Surrey, GU9 7PT
England

Ashgate Publishing Company
Suite 420
101 Cherry Street
Burlington
VT 05401-4405
USA

www.ashgate.com

British Library Cataloguing in Publication Data
Resilience engineering in practice : a guidebook. --
(Ashgate studies in resilience engineering)
1. Human engineering. 2. Industrial safety. 3. Risk perception.
I. Series II. Hollnagel, Erik, 1941-
620.8'2-dc22

ISBN: 978-1-4094-1035-5 (hbk)
ISBN: 978-1-4094-1036-2 (ebk) I

Library of Congress Cataloging-in-Publication Data
Resilience engineering in practice : a guidebook / by Erik Hollnagel ... [et al.].
p. cm. -- (Ashgate studies in resilience engineering)
Includes index.
ISBN 978-1-4094-1035-5 (alk. paper) -- ISBN 978-1-4094-1036-2 1. Reliability (Engineering) 2. Fault tolerance (Engineering) I. Hollnagel, Erik, 1941-
TA169.R466 2010
620'.00452--dc22

2010032902

FSC
www.fsc.org
MIX
Paper from responsible sources
FSC® C018575

Printed and bound in Great Britain by the
MPG Books Group, UK

Contents

This page has been left blank intentionally

List of Figures

List of Tables

This page has been left blank intentionally

List of Contributors

Karina Aase is professor in safety at the University of Stavanger, Faculty of Social Sciences. She is an engineering graduate from the Norwegian University of Science and Technology and holds a PhD in work environment and safety with a thesis on experience transfer within the petroleum sector. Aase has conducted research within safety and organisational learning within different sectors such as petroleum, transport, and healthcare during the last ten years. Her research topics are high-risk organisations, safety, learning, human–technology–organisation–society interfaces, change processes and socio-technical systems. Currently Aase manages research projects within aviation and patient safety.

Rene Amalberti is professor of medicine, physiology, and ergonomics. He joined the Air Force in 1977 and retired with the rank of general in February 2008. He is now Senior Advisor, Patient Safety, at the Haute Autorité de Santé, and risk manager in a medical insurance company. From 1982–1992, he was involved in several major European research programs on human error and risk management, and pioneered the concepts of pilot's assistant, ecological safety, and Crew Resource Management. From 1992–2000 he studied risk management in the nuclear and oil industry, professional fishing, and public ground transportation. In the late 1990s, he undertook research in the medical field on medical errors, patient safety, system approach, and resilience.

Ion Berechet is an engineer and physicist. In 2002, he joined Air France Consulting. In 2003, Ion created and developed SISPIA (Intelligent System for Surveillance Inaccessible Parameters in Aeronautical process for risk management and performance). He is project leader for safety systems, risk analysis and control optimisation for aeronautical ground operations for Air France – KLM group. Ion has put in place a methodology and mathematical analysis of ASR for the definition of safety management systems for the risk of fatigue for short couriers with DGAC.

Johan Bergström has a MSc in Risk Management Engineering and is a PhD student in the field of Systems Safety at Lund University in Sweden. He is currently working on a research project concerning organisational resilience in escalating situations.

Matthieu Branlat is a PhD student at the Ohio State University, Columbus, OH. He received a MSc in Ergonomics from *Conservatoire National des Arts et Métiers*, Paris, France in 2006. His research interests include resilience engineering, system safety, decision making, human expertise and collaborative work. Recent or on-going projects have been conducted in domains such as disaster management, urban fire-fighting, intelligence analysis and healthcare.

Philippe Cabon is currently assistant professor in Human Factors at Paris Descartes University (France). He received his MSc in Psychology and Neuroscience and a PhD from Paris Descartes University. His main research interests include fatigue, workload and stress in Civil Aviation, Air Traffic Control and industry. For the last 15 years he has worked extensively on fatigue and sleep of aircrews in long and short-range flights in cooperation with the French Civil Aviation Authorities (DGAC), Airbus and several airlines. He is currently working on the implementation of Fatigue Risk Management System in civil aviation.

Lucie Cuvelier is a safety engineer and PhD candidate in Ergonomics at the Center for Research on Work and Development – Ergonomics Lab – within the *Conservatoire National des Arts et Métiers* (CNAM, Paris). Her research is in the field of patient safety and aims to develop new approaches of reliability and industrial risk-management. Its central theme concerns the identification and development of organisational pre-conditions that enable the anesthesiologists to cope with daily variabilities as well as with unforeseen situations.

Nicklas Dahlström is a Human Factors researcher and Crew Resource Management instructor at Lund University School of Aviation, Sweden, as well as a Human Factors specialist at Emirates Airlines in Dubai. He has been an officer in the Swedish Air Force, has a BSc in Meteorology from Stockholm University and a PhD in Technology from Lund University. His main research areas are mental workload, primarily with aviation but also in maritime and railway transportation, nuclear and chemical industry, medicine and rescue services.

Stephane Deharvengt is a civil aviation engineer with a PhD in Ergonomics. He spent more than 15 years with the French National Safety Authority, DSAC, as an expert in safety risk-management related to organisational and human factors for rulemaking, cockpit and cabin certification of aircraft, and implementation of safety management systems for airlines and maintenance and repair organisations. He is now deputy head of safety, quality and security for the French Air Navigation Service Provider, DSNA.

Sidney Dekker is professor of Systems Safety Factors & Aviation Safety at Lund University in Sweden. He gained his PhD in Cognitive Systems Engineering at the Ohio State University. His books include *The Field Guide to Human Error Investigations* and *Ten Questions about Human Error*. His latest book is *Just Culture: Balancing Safety and Accountability*. Sidney Dekker is also an airline pilot, flying the Boeing 737NG part-time, as well as a member of the Union of Concerned Scientists.

Pierre Falzon is professor of Ergonomics and Neurosciences of work at the *Conservatoire National des Arts et Métiers* (CNAM) in Paris, France. He heads the Ergonomics Lab, part of the Centre de Recherches sur la Travail et le Développement. His present research interests concern processes of competence development and knowledge construction and human and organisational reliability. On a more general level, he is interested in epistemological issues related to the practice of ergonomics and to ergonomics as a discipline. Pierre Falzon is past President of the International Ergonomics Association.

Pedro Ferreira is concluding a university placement program at Network Rail in London, to obtain a PhD from the University of Nottingham, UK. His research is based on the subject of resilience in the planning of rail-engineering work. Pedro holds a MSc from the Faculty of Human Kinetics, Technical University of Lisbon. He developed his thesis on the subject of human reliability and the occurrence of violations railway traffic control. He has worked on several industrial sectors as a human factors consultant and lectured Health and Safety courses at the Technical University of Lisbon.

Jean-Yves Grau started his career as a senior scientist in the French Military Institute of Aerospace Medicine (IMASSA). After completing his studies in Medicine and specialisation in aviation medicine, he specialised in aviation psychology. He was chief of IMASSA flight safety office and has, since 2001, managed a private consulting firm in Human Reliability in transport and risky industries. He is currently working within the STARE consortium in order to design methodology for managing 'fatigue' risk in air operations. He is a lecturer in Human Factors, Flight Safety and Risk Management.

Erik Hollnagel is professor and Industrial Safety Chair at MINES ParisTech (France) and Visiting Professor at the Norwegian University of Science and Technology (NTNU) in Trondheim (Norway). He has worked at universities and research centres, and in industry in several countries and with problems from many domains. His professional interests include industrial safety, resilience engineering, accident investigation, cognitive systems engineering and cognitive ergonomics. He has published widely and is the author/editor of 17 books, including four books on resilience engineering. The latest title from Ashgate is *The ETTO Principle: Why Things That Go Right, Sometimes Go Wrong*.

Daniel H. Hummerdal is a psychologist and currently studying for a PhD on the cost of safety at MINES ParisTech, France. After completing his commercial pilot training he worked as an instructor at flying schools in both Sweden and the US. He has also worked for five years as an accident investigator for the Swedish Civil Aviation Authority. As a guest lecturer at Lund University he carried out research on patient safety and developed teaching material for Crew Resource Management courses. Before enrolling as a PhD student he worked as a human factors consultant with Dédale, France, evaluating the EU's flight safety research program – HILAS.

Tom Kontogiannis is associate professor in Human Factors and Industrial Safety at the Department of Production Engineering and Management, Technical University of Crete. He holds degrees in mechanical engineering, industrial safety and cognitive ergonomics. While working in the UK he took part in more than 30 industrial projects concerning aspects of human reliability, modeling of human performance under stress, accident investigation and design of decision-support systems. Since 1997, he has headed a small group at the Cognitive Ergonomics & Industrial safety (CEIS) Lab that focuses on human-performance modeling, workflow modeling and industrial safety. He has published over 30 peer-reviewed journal papers and a book, available at http://www.dpem.tuc.gr.

Elizabeth Lay is manager of the Field Service Operational Risk Management Group at Siemens Energy. She graduated from the University of North Carolina at Charlotte with a BSc in Mechanical Engineering and is currently pursuing a Certificate of Cognitive Sciences at the University of Central Florida. She has held various roles in the business of servicing power plants leading to her current role focused on risk management and loss control. Outside of work, Elizabeth is a photographer and avid world traveler, visiting at least one new country every year.

Nicolas P. Maille is a research scientist in the human factor team of ONERA. He held a PhD in artificial intelligence from SUPAERO in 1999. He has conducted research to improve experience feedback methodologies and tools in the aeronautical field with major aeronautical actors including NASA, Airbus, French Air Forces and commercial airlines. He has a primary interest in the analysis of pilot activity through both crew reports and digital flight data.

Stathis Malakis holds a mathematics degree, an MSc in Air Transport Management and a PhD in Cognitive Systems Engineering. Since 1999 he has worked for the Hellenic Civil

Aviation Authority as an Air Traffic Controller holding Tower, Approach, Terminal Approach Radar and Instructor ratings. In 2005, he joined the department of Production Engineering and Management at Technical University of Crete to study for a PhD under the supervision of Associate Professor Tom Kontogiannis. He is actively involved in ATC training and safety issues both at professional and research levels.

Régis Mollard is professor in Ergonomics at University Paris City – Descartes. He holds a PhD in 3D Surface Anthropometry and Human Modeling. Responsible of Master Diploma in Ergonomics and University Diploma in Human Factors in aeronautics. He is a member of the World Engineering Anthropometry Resource (WEAR) group, the European Committee for Aircrew Scheduling and Safety (ECASS), the French Society of Biomechanics and Psychophysiology in Ergonomics (PIE), a technical committee of the International Ergonomics Association (IEA).

Anne Sophie Nyssen is doctor of Work Psychology and Professor of Cognitive Ergonomics at the University of Liege, Belgium. Her primary research focused on human errors in anesthesia but has been extended to other work situations such as process control and transportation. More recently, her research focuses on robotic surgery to assess the impact of the use of robotic surgery using 3D on knowledge, strategies and errors. She is the author and co-author of many international peer-reviewed publications.

Jean Pariès graduated from the French National School of Civil Aviation as an engineer, then joined the Direction Générale de l'Aviation Civile) (DGAC), dealing with air safety regulations. He was a member of the ICAO Human Factors and Flight Safety Study Group since its creation in 1988. In 1990, he joined the Bureau Enquêtes Accident as Deputy Head, and Head of Investigations. In 1994, he left the BEA to found Dédale SAS. From 2000 to 2004, he was also a Research Associate Director with the Centre National de Recherche Scientifique (CNRS) and worked on hazards associated with failures in high risk activities, and their control by individuals, teams and organisations. He is the author of numerous papers, book chapters, and communications on Human and Organisational Factors in safety.

Alberto Pasquini achieved his MSc in Engineering from the University of Rome La Sapienza in 1978. He has been with the licensing authority for nuclear power plants of the Italian Research Agency ENEA for several years. He is now with the Robotic Division of the same agency and R&D Director at Deep Blue, an Italian consultancy and research company operating in the domain of air transportation. His research interests are in system safety, human machine interactions and validation. He has more than 60 publications in his area of professional activity in scientific journals, books, and conference proceedings.

Kurt Petersen is professor of Risk Analysis and Management at Lund University in Sweden. He gained his PhD at the Technical University of Denmark. He has been at Risø National Laboratory and Danish Transport Research Institute, both in Denmark, prior to the professorship at Lund University. His has been working with safety assessment at nuclear, off-shore, chemical systems, and large infrastructure system like natural gas, electric power, and railway. His main research interests are methods for risk analysis, risk management, risk and vulnerability analysis, and accidents investigation.

Simone Pozzi works as Human Factors and Safety R&D expert in Deep Blue Consultancy and Research (Rome). He also lectures Interaction Design at the Department of Architecture and Urban Planning, University of Sassari at Alghero. He achieved his MSc (2001) and PhD (2006), with specialisation in human–computer interaction. He has participated in many international R&D projects in the area of Human Factors and Safety. He regularly publishes in peer-reviewed international academic journals, books, and conference proceedings.

Brendan Ryan has a PhD from the University of Nottingham, UK. He has particular research interests in self reporting and interviewing in accident or other work situations, and in the identification, assessment and management of risk in a range of industrial contexts. Brendan has worked in industry as an accident investigator, in safety and environmental health consultancy and government inspectorate roles. He has a recent and substantial body of research work supporting engineering and sustainability projects in the rail industry. He is a Chartered Member of the Institution of Occupational Safety and Health.

Luca Save works as Human Factors and Safety R&D expert in Deep Blue Consultancy and Research (Rome). He has ten years of experience in the area of safety critical systems, including air traffic management, railway transport and clinical risk management. In 2003 he achieved a PhD in human–computer interaction, with a special focus on cognitive ergonomics. He has participated in research initiatives founded by the Italian Railway Organisation and in many international R&D projects in the area of Human Factors and Safety. He has worked as technical consultant for the investigation of railway accidents for the Italian Public Prosecutor office of Bologna in 2005 and 2006. He has a number of publications in peer-reviewed international academic journals and in international conferences.

Sarah Sharples is Associate Professor in Human Factors and Head of the Human Factors Research Group in the School of Mechanical, Materials and Manufacturing Engineering at the University of Nottingham. She is member of the EPSRC funded Rail Research UK project (RRUK2) in which she is leading work on analysing the impact of automation in railway control. She has also worked on and managed a number of EU projects in several industrial contexts, including automotive and aerospace design and manufacture. Her main areas of interest and expertise are transport human factors, human-computer interaction, cognitive ergonomics and development of quantitative and qualitative research methodologies for examination of interaction with innovative technologies.

John Stoop graduated in 1976 as an aerospace engineer at Delft University of Technology (DUT). He is one of the founders of the Safety Science Group at DUT. He works part-time at the Faculty of Aerospace Engineering at DUT and is guest professor at Lund University in Sweden. He is involved in developing a safety oriented methodology for in-depth accident investigation, forensic engineering and the integration of safety into the design process. He is Affiliated Member of the International Society of Air Safety Investigators ISASI, member of the Board of Dutch Association of Road Victims VVS, and founder and managing director of Kindunos Safety Consultancy Ltd, established in 1990.

Mark-Alexander Sujan is Assistant Professor in Patient Safety at Warwick Medical School, University of Warwick. He holds degrees in Computer Science (MSc), Human Factors (PhD), and Philosophy, Psychology and Political Science (BA). His current research interest is on approaches to enhance the reliability of healthcare systems. He leads the Masters module Patient Safety at Warwick and delivers the Reliability in Healthcare course for the NHS Institute for Innovation and Improvement. Mark is also a senior Fellow of the NHS Improvement Faculty.

Gunilla A. Sundström is a global professional and consultant for Financial Services, Outsourcing/Offshoring and R&D. She has held leadership positions in a variety of industries including R & D, Financial Services, Global Sourcing and currently holds the position as the Head of Global Sourcing with Deutsche Bank. Dr Sundström has published more than 60 papers; holds two US Patents and has been awarded IEEE-Systems, Man and Cybernetics' outstanding contributions award. Dr. Sundström holds a DPhil from University of Mannheim, Germany.

Berit Berg Tjørhom is completing her PhD work on risk management and societal safety at the University of Stavanger, Faculty of Social Sciences. She holds a master in societal safety with a thesis on safety culture within aviation. Her PhD thesis is entitled *Exploring Risk Governance in a Global Transport System*, a thesis that contributes to an understanding of risk governance within the Norwegian aviation transport system. Her research focuses on risk management in a human–technology–organisation–society perspective related to interdependencies, interfaces and complexity.

John Wilson is professor of Human Factors at the University of Nottingham, and is also Principal Ergonomist at Network Rail. John is a Chartered Psychologist and a Chartered Engineer, a Fellow of the Ergonomics Society, and was awarded the Ergonomics Society Sir Frederic Bartlett Medal in 1995 and the Distinguished Overseas Colleague Award of the Human Factors and Ergonomics Society in 2008. His research contributions to rail human factors include the development of performance assessment approaches for workload and situation awareness, human factors of track access and protection arrangements, and the understanding of human factors risk generally.

David Woods is professor at the Ohio State University in the Department of Integrated Systems Engineering and Director of the university's initiative on Complexity in Natural, Social and Engineered Systems. He has developed and advanced the foundations and practice of Cognitive Systems Engineering since its beginning. He has studied team work between people and automation in anesthesiology, aviation, space mission operations, and health care. A past President of the Human Factors and Ergonomic Society he has served on US National Academy of Science and other advisory committees. He was a board member of the National Patient Safety Foundation during its startup, Associate Director of the Midwest Center for Inquiry on Patient Safety of the Veterans Health Administration, and advisor to the Columbia Accident Investigation Board.

John Wreathall is a specialist in systems engineering methods with particular emphasis on human and organisational performance as it relates to safety, reliability and quality. He has pioneered the development and application of human-performance analysis methods, both quantitative and qualitative, for application in the medical, transportation, nuclear, and aerospace communities. He has participated in the development of the understanding of human errors and the circumstances that lead to their occurrence, including chairing and presenting at international conferences on this subject. He was the invited keynote speaker on the subject of resilience engineering at the First Mercosur (South American Free

Trade Association) Conference on Safety and Security in Work and its Environment, Porto Alegre (Brazil) in 2004.

Kyla Zimmermann (née Steele) is a human factors analyst for the Transportation Safety Board of Canada. Her career started in aerospace engineering, in the testing and development of aircraft and aero engines, as well as trouble-shooting airline technical issues. Kyla moved to Europe in search of adventure and a career transition to human factors. Her MSc research at Linköping University examined aviation incident reporting systems in Sweden. In France Kyla pursued doctoral studies examining safety models and perspectives, in particular the differences in perspectives across occupations and societal cultures in the aviation industry.

Reviews for *Resilience Engineering in Practice*

Although risk management has brought greater safety to socio-technical systems, a new approach is still strongly needed. Erik Hollnagel's excellent book offers the right approach; that resilient behaviour by people leads to stable systems. Those searching for a more profound understanding of system safety must read this book as it is a practical guide to this new approach.

Akinori Komatsubara, Waseda University, Japan

With crises abounding, the concept of resilience is more relevant than ever. Manifold examples from a variety of high-risk industries provide insights into the four basic requirements for resilience: responding, monitoring, anticipating, and learning. Tools are presented that support the assessment of these requirements as well as their promotion, be it by training emergency management, handling fatigue of system operators, supporting preventive maintenance, providing better rules for managing conflicting goals, or improving incident reporting. The book, by Erik Hollnagel and his colleagues, will be a great resource for system designers and decision-makers in organizations in their endeavours to keep the uncertainties and complexities of our world at bay.

Gudela Grote, ETH Zürich, Switzerland

'Be prepared to be unprepared.' How do you do that? By absorbing the evocative data, nuanced terminology, sustained guidance, and broad applications summarized here. Resilience is about more than engineering as becomes clear in these descriptions of the actual, critical, potential, and factual events that unfold when 'disturbances fall outside the operational envelope.' Resilience engineering is a hot topic. Here is the one book that shows you why!

Karl E. Weick, University of Michigan, USA

Prologue: The Scope of Resilience Engineering

Erik Hollnagel

The focus for safety efforts is usually, and traditionally, the unwanted outcomes, injuries, and losses that are the result of adverse events. This matches the common understanding of safety as 'freedom from unacceptable risk.' Resilience Engineering, however, defines safety as the ability to succeed under varying conditions. It is a consequence of this definition that it is equally important to study things that go right as things that go wrong. For Resilience Engineering, the understanding of the normal functioning of a socio-technical system is the necessary and sufficient basis for understanding how it fails. And it is both easier and more effective to increase safety by improving the number of things that go right, than by reducing the number of things that go wrong. The definition of resilience can be made more concrete by pointing to four abilities that are necessary for a system to be resilient. These are the ability to respond to events, to monitor ongoing developments, to anticipate future threats and opportunities, and to learn from past failures and successes alike. The engineering of resilience comprises the ways in which these four capabilities can be established and managed.

Introduction

In the world of safety, comprising issues such as accident investigation, risk assessment, safety management, and safety culture, the focus has traditionally been on that which has gone wrong or could go wrong. This is illustrated by the commonly used definition of safety as 'freedom from unacceptable risk.' The focus on what could go wrong obviously makes practical sense, since clearly it is important for every enterprise to understand both what has gone wrong and what may go wrong, in order to develop measures either to prevent it from happening (again) or to protect against the outcomes.

This line of thinking is well illustrated by the traditional risk matrix, an example of which is shown in Figure P.1. The risk matrix characterises the risk level of possible outcomes by considering both their probability, that is, the likelihood that they will happen, and the severity of the consequences, that is, the magnitude of the possible consequences.

The risk matrix, however, only looks at things that can go wrong. Yet if we consider the possible outcomes of something (an event, a function, or a process), it is clear that things can go right as well as wrong. It is furthermore reasonable to expect that things normally will go right, that they will turn out as planned or intended, and that it is unusual for things to go wrong. We are therefore unpleasantly surprised when it happens. In view of this, it seems reasonable to propose that a description of possible outcomes should go beyond the traditional risk matrix and extend the 'consequence' dimension to include both positive (wanted) and negative (unwanted) outcomes. This can be shown as in Figure P.2.

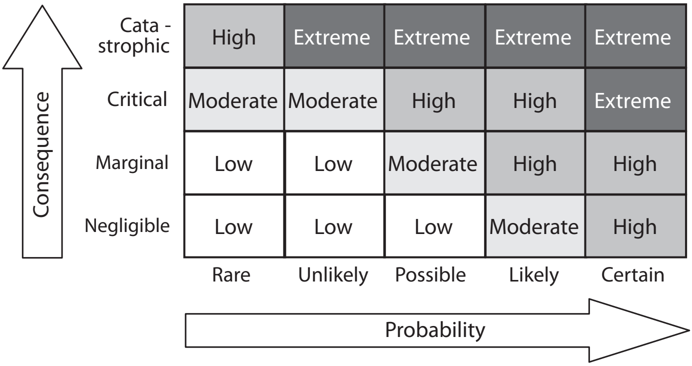
Consequence
Probability

Cata - strophic	High	Extreme	Extreme	Extreme	Extreme
Critical	Moderate	Moderate	High	High	Extreme
Marginal	Low	Low	Moderate	High	High
Negligible	Low	Low	Low	Moderate	High
	Rare	Unlikely	Possible	Likely	Certain

Figure P.1 A traditional risk matrix

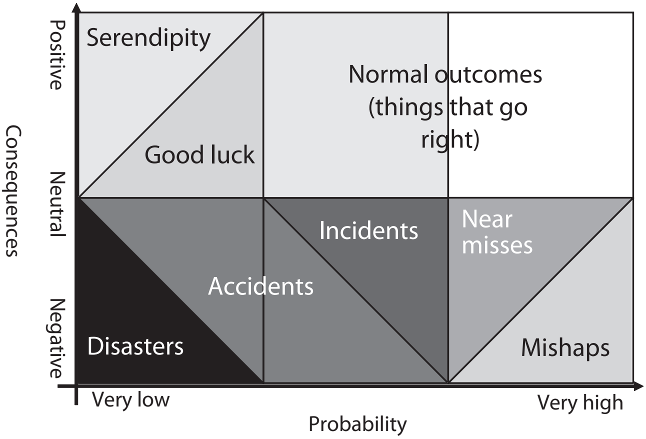

Figure P.2 Range of outcomes

Safety efforts have traditionally focused on unwanted or negative outcomes, and have furthermore been limited to outcomes with a relatively low probability such as incidents and accidents. (Unwanted negative outcomes with high probability, for example, mishaps, will normally have been eliminated, since otherwise the system would be unable to maintain its required functioning.) If we for a moment assume that there is a simple causal relation between events and outcomes, it becomes possible to characterise several characteristic subsets of outcomes as follows:

1. Positive outcomes that have a high probability. This subset represents the successes or 'normal' actions, that is, the things that not only go right but that also are intended and expected to go right—in other words, normal work or normal functioning. Indeed, if normal work either did not result in wanted outcomes, or was not highly predictable, something would be seriously wrong.
2. Positive outcomes that have a low probability. This subset represents the 'good' things that can happen, but that happen unexpectedly. There is no commonly recognised terminology for these, but terms such as *serendipity* or even *good fortune* represent at least some of them.

3. Negative or unwanted outcomes that have a low probability, that is, things that go wrong and which are unexpected—although not unimaginable. This is the subset of outcomes that usually is associated with safety—or rather, with a lack of safety—particularly outcomes that are serious (in terms of causing significant losses) and that are hard to predict. This subset includes the commonly used categories of incidents and accidents. It also includes disasters, although these rarely are covered by industrial safety.
4. Negative or unwanted outcomes that have a high probability. This basically means outcomes that realistically must be expected to happen frequently or even regularly. In practice most of these outcomes have only minor negative consequences, because they otherwise would have been eliminated (cf., the As Low As Reasonably Practicable (ALARP) principle). They are commonly described as near misses or 'almost accidents,' or as unsafe actions. Near misses are usually benign but may lead to serious negative consequences. Another subset is the mishaps, that is, 'near misses' with serious outcomes. Predictable events that may result in serious unwanted outcomes can, however, normally be are assumed to have been eliminated.

A more condensed description of the four sets of outcomes is shown Table P.1.

Table P.1 The sets of possible outcomes

	Things that go right (wanted outcomes)	**Things that go wrong (unwanted outcomes)**
Outcomes with high predictability	This is the set of outcomes that represent the normal functioning of a safe system. Ought to be governed by an As High As Reasonably Practicable (AHARP) principle.	The serious outcomes in this set are normally eliminated; the minor unwanted outcomes are usually tolerated, as described by the ALARP principle.
Outcomes with low predictability	These outcomes are not normally considered in system management, but should obviously be facilitated as far as possible. They are gratefully accepted if and when they occur.	These outcomes are the focus of traditional safety efforts. They are the subject of risk assessment, prevention, and protection. Many efforts are made to calculate how 'unexpected' they are, hence transfer them to the set above.

As already mentioned, safety efforts have usually focused on outcomes that are both unwanted (i.e., with significant negative consequences) and unexpected or difficult to predict, corresponding to the categories of accidents and incidents in Figure P.2 or the high to extreme risks in Figure P.1. The common understanding is that safety can be achieved if accidents, incidents (and mishaps) either can be prevented or if their number (or frequency) can be reduced. Disasters must, of course, not be neglected although their predictability usually is so low that it is difficult to do much to prepare for them. (In relation to the terminology proposed by Westrum (2006), disasters can be seen as irregular threats or even improbable events.) Since the 1980s, the safety focus has occasionally been extended from incidents and mishaps to include near-misses also. But the practical problem is that there are so many near misses, that they happen so frequently, and that the consequences usually are negligible so that it is not considered cost-effective to do much about them.

More importantly, the traditional approaches to safety usually disregard what lies 'above' the middle of Figure P.2, that is, the ways in which things can go right. This is due to the unspoken assumption that we can best learn about things that go wrong by studying only things that go wrong. It is nice when things go right, but there is no need to pay much attention to them precisely because they go right. It is also due to the fact that as we get used to something, we tend not to notice it any longer. (The technical term for this is habituation, which denotes the psychological process in humans that leads to a decrease in response to a stimulus after repeated exposure over a specified duration of time.)

Resilience Engineering, however, takes a different position. Resilience Engineering sees the 'things that go wrong' as the flip side of the 'things that go right,' and therefore assumes that they are a result of the same underlying processes. In consequence of that, 'things that go right' and 'things that go wrong' should be explained in basically the same way. It therefore makes as much sense to try to understand why things go right as to understand why they go wrong. In fact, it makes more sense because there are many more things that go right than things that go wrong, the ratio depending on how (im)probable an accident is considered to be. If, for instance, the probability of failure is 10E–4 (meaning 10–4 or 1/10.000), then humans are usually blamed for 80–90 percent of the one case out of 10.000 when things go wrong. By the same 'logic,' humans should be praised for a similar 80–90 percent of the 9.999 cases where nothing goes wrong. (In both cases humans should actually be seen as accountable for the full 100 percent, since it would otherwise be necessary to postulate some *deus ex machina* to account for the remaining 10–20 percent.) Resilience Engineering proposes that we should try to understand a system's performance in general, rather than limit ourselves to the things that go wrong, that is, try to understand all the outcomes shown in Figure P.2 rather than only the negative ones—with the possible exception of 'good luck.'

Both Figures P.1 and P.2 use the probability of an outcome as a descriptive dimension, but neither considers the frequency of outcomes. While probability and frequency are closely linked, they do not mean the same, and for the safety of everyday work, the frequency of outcomes is perhaps the more important. Following the argument made above, if things go wrong one time out of every 10.000, then things go right the remaining 9.999 times. This is illustrated in Figure P.3, where a third dimension, representing frequency, has been added to the diagram shown in Figure P.2.

As Figure P.3 tries to illustrate, there are many more things that go right than things that go wrong. Even for ultra-performing systems, the ratio is around 1:1.000 (Amalberti, 2006). For ultra-safe systems it may be 1:1.000.000 or even lower, meaning that the number of normal outcomes is at least six orders of magnitude larger than the number of failures. It is the set of normal outcomes that rightly ought to represent the safe performance of a system or process, just as the set of accidents and incidents represent unsafe performance. It may therefore be said that safety efforts, almost paradoxically, have focused on unsafe functioning rather than on safe functioning. This may, as noted above, be due to the psychological fact that safety is nearly invisible while a lack of safety is highly visible. We notice that which is unusual while we become habituated to that which is usual. Resilience Engineering recognises this paradox and argues that safety should deal with

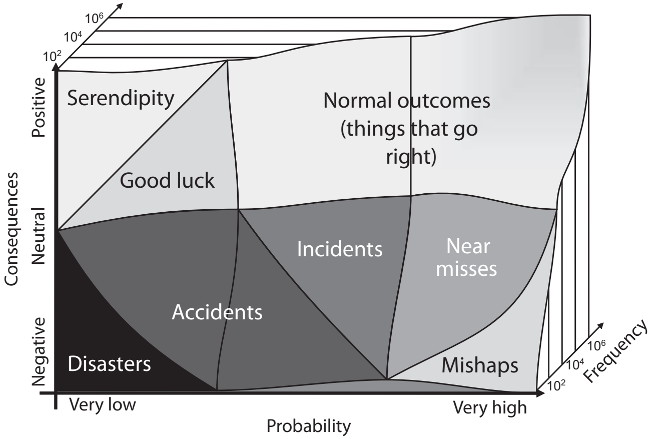

Figure P.3 The frequency of various outcomes

safe performance as well as unsafe performance—with things that go right as well as things that go wrong.

According to this line of reasoning, Resilience Engineering is not a simple replacement for safety (management). Safety (management) has traditionally, and with good reason, focused on a subset of the possible events and outcomes. This was in many ways sufficient as long as systems and processes were manageable, or tractable, so that normal functioning could be ensured by limiting or constraining performance variability (cf., Hollnagel, 2009). The developments in socio-technical systems during the last 20 years or so have, however, created an increasing number of systems and processes that are intractable, and where performance variability consequently is a necessity and an asset rather than a liability. Resilience Engineering argues that it is necessary to look at success as well as at failures precisely in order to understand failures or why things wrong. The argument is that there are no special 'error producing' processes that magically begin to work when an accident is going to happen, but which otherwise lie dormant. On the contrary, there are no fundamental differences between performance that leads to failures and performance that leads to successes. We are therefore best served by trying to understand performance in general, regardless of whether we focus on individual, collective, or organisational performance.

The difference between 'classical' safety management and Resilience Engineering is demonstrated by the differences between the definitions. A common definition of safety was mentioned above. Resilience can in the same manner be defined as:

> The intrinsic ability of a system to adjust its functioning prior to, during, or following changes and disturbances, so that it can sustain required operations under both expected and unexpected conditions.

This definition obviously comprises the classical definition of safety, since 'the ability to sustain required operations' is tantamount to the 'freedom from unacceptable risks.' But the definition of resilience also makes clear that safety cannot be seen independently of the core process (or business) of the system, hence the emphasis on the ability to *function* under 'both expected and unexpected conditions' rather than just to avoid failures. It is this ability that makes the system both safe and efficient, and Resilience Engineering deals with both.

The difference between the two views is illustrated by Figure P.4, which uses a balance to show two different ways to improve safety. One is to reduce the number of things that go wrong, which obviously will tip the scale in favor of safety. The other is to increase the number of things that go right, which will achieve the same effect, but which at the same time will contribute to productivity and the core business processes. Resilience Engineering favors the second approach. The goal of Resilience Engineering is to increase the number of things that go right rather than to reduce the number of things that go wrong, noting that the latter will be a consequence of the former.

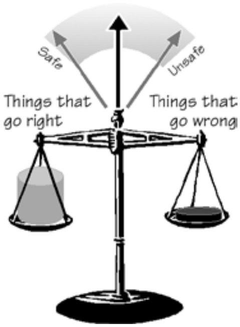

Figure P.4 A resilience engineering view of safety

The Four Cornerstones of Resilience

If we define resilience as proposed above, the goal of Resilience Engineering becomes how to bring about resilience in a system. The key term of this definition is the system's ability to *adjust* its functioning. This working definition of resilience can be made more detailed by noticing that it implies four main factors, each representing an essential system ability. The four factors, or four essential abilities are (cf., Figure P.5):

- Knowing what to *do*, that is, how to respond to regular and irregular disruptions and disturbances either by implementing a prepared set of responses or by adjusting normal functioning. This is the ability to address the *actual*.
- Knowing what to look for, that is, how to *monitor* that which is or can become a threat in the near term. The monitoring must cover both that which happens in the environment and that which happens in the system itself, that is, its own performance. This is the ability to address the *critical*.
- Knowing what to *expect*, that is, how to anticipate developments, threats, and opportunities further into the future, such as potential changes, disruptions, pressures, and their consequences. This is the ability to address the *potential*.
- Knowing what *has happened*, that is, how to learn from experience, in particular how to learn the right lessons from the right experience—successes as well as failures. This is the ability to address the *factual*.

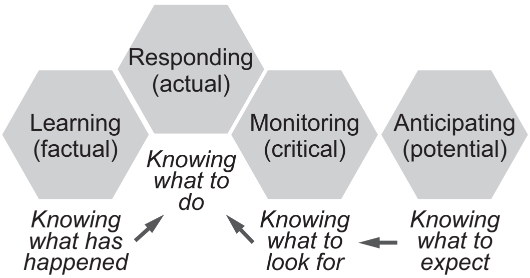

Figure P.5 The four cornerstones of resilience

If the resilience of a system is defined by the abilities to respond to the actual, to monitor the critical, to anticipate the potential, and to learn from the factual, an obvious question is how this can be brought about? This is really the question of how resilience can be *engineered* or the question of what Resilience Engineering is in practice. A detailed answer, or rather detailed answers, can be developed by considering each of the four factors in a more operational perspective. This quickly leads to a number of issues that can serve as the starting point for more concrete measures (cf., Epilogue). The focus on the issues arising from each of the four factors provides a way to think about Resilience Engineering in a practical manner. Starting from the level of the system as a whole this soon leads to the development of operational details and specific steps to be taken on a concrete level. This can, however, only be done by referring to a specific domain or field of activity, or even to a specific organisation at a certain time. Much of that may obviously make use of existing methods and techniques, although seen from a resilience perspective and in some cases be supplemented by new methods and techniques. For any given domain or organisation it will also be necessary to determine the relative weight or importance of the four main abilities, that is, how much of each is needed. The right proportion cannot be determined analytically, but must be based on expert knowledge of the system under considerations and with due consideration of the characteristics of the core business. Yet the minimum requirement is that none of the four can be left out if a system wants to call itself resilient.

Reading Guide

The chapters of this book have been organised in four main sections that correspond to the four main resilience abilities. While this organisation serves to emphasise how each ability can be considered in more detail, the chapters also make clear that Resilience Engineering cannot work by focusing on each of the four abilities in isolation. The four abilities depend on each other, and it is necessary to acknowledge and understand the dependencies or couplings among them in order successfully to ‘engineer’ resilience. Corresponding to the ‘new’ definition of safety as the ability to succeed under varying conditions, the four abilities represent functions that can be improved, hence something that grows as the safety of a system gets better. Taken together, strengthening the abilities to respond, to monitor, to anticipate, and to learn is the best way to ensure that more things go right and that fewer things go wrong.

This page has been left blank intentionally

PART I
Dealing with the Actual

This page has been left blank intentionally

Chapter 1
Resilience and the Ability to Respond

Jean Pariès

Resilience in 'Real Time'

In this first section, the emphasis is on the ability of an organisation or a system to 'deal with the actual,' that is, to respond to the demands of the current situation—a disrupting situation. At the 'sharp end' of the system, 'responding to the situation' includes assessing the situation, knowing what to respond to, finding or deciding what to do, and when to do it. The readiness to respond mainly relies on two strategies. The first—and proactive one—is to anticipate the potential disruptive situations and predefine ready-for-use solutions (e.g., abnormal or emergency procedures, specific reaction skills, crisis response plans, and so on). The second—and reactive one—is to generate, create, invent, or derive ad hoc solutions.

To put it differently, this section is about 'real time' resilience, but from a synchronic as well as diachronic perspective. Indeed, at the 'blunt end' of the system (i.e., the domain of designers, managers or trainers), issues related to 'real time resilience' include how to ensure that the required resources (people, competence, equipment) are available or can be established in time. Hence a more complete question would be how to establish (now) and maintain (tomorrow) a readiness to respond (at any time in the future). While the three chapters in this section focus on the 'readiness to respond,' they also tackle some issues related to establishing and maintaining that readiness. Each of them presents a practical case taken from a specific domain:

commercial aviation, anaesthesia, and rescue services. Beyond their obvious difference of perspective and domain, they share similar underlying theoretical questions. The first of these is the relationship between resilience and anticipation, which also runs throughout other chapters of this book.

Readiness and Anticipation

At first glance, the role of anticipation is both obvious and simple: things go better when they have been anticipated. In *Lessons from the Hudson* (Chapter 2 this volume), the author revisits the successful ditching of US Airways Flight 1549 into the Hudson River in 2009. He describes a 'defence in depth' strategy against the anticipated bird hazard, engineered into the aviation system: the first line of defense is to minimise the frequency of bird strikes, the second is to assure the ability of the aircraft and its engines to withstand hitting some birds without damage, and so on. The last line of defense is the ability of the crew-aircraft system to land on unprepared terrain or ditch with minimum damage after a total loss of power, and to evacuate passengers safely. Clearly, the fact that the aviation system had anticipated a total engine failure greatly contributed to survivability in the Hudson River event. The aircraft design allowed the crew to keep some control of the flight path, dual engine failure and ditching procedures were available, and the crew had been trained in emergency evacuation. However, while fully anticipated at the aviation system level, a dual engine failure is an extremely remote event at the scale of a pilot's professional life, and its occurrence came as a total surprise for the crew. This highlights that anticipation is not something that is uniformly distributed throughout a large system. The global system may anticipate occurrences that are too rare to be even thought of at local scales, while local operators will anticipate situations that are much too detailed to be tackled at a larger scale. This raises the issue of the coupling between the different levels of organisation within a system, what Woods (2006a) calls 'cross-scale interactions.' We will come back to this later.

In Chapter 3, Cuvelier and Falzon discuss how anaesthetists manage critical situations. They show that practitioners anticipate a specific range of 'potential variability' before each operation, prepare the necessary responses and resources, and feel that they are in control as long as the 'unexpected' events stay within the boundaries of this area. When an event falls outside this area (mostly because of equipment failures or cooperation problems with the surgical team), the first challenge is to recognise it as such and anaesthetists found themselves with problems in most cases either having to identify and understand the situation or implement appropriate responses under very tight time constraints. Consequently, the authors differentiate 'unexpected' events according to the nature of the related surprise. They classify events as '*potential situations*' when they had been envisaged by the anaesthetists before the operation (while of course not expected to happen at this time and place), and as '*unthought-of situations*' when they had not been envisaged at all, similar to the 'unprecedented events' of Westrum's (2006) classification. In nine of the thirteen so-called unthought-of situations, the anaesthetists *called on colleagues for help*. The authors see these decisions to call on additional resources (particularly to call on colleagues) as the observable sign of a shift from controlled to crisis situations. They argue that resilience lies in the operator's ability, not only to detect, but also to accept—and literally '*decide*'—that the system has breached the boundaries of potential variability. This is a key point; a crucial condition for maintaining and/or recovering control is indeed the ability to detect, recognise and accept that the situation is beyond what had been imagined by the operators on the basis of their experience envelope. Many accidents can be understood as the result of a failure to recognise/accept an excursion of the real situation outside the range of anticipated variations, leading to a continuation of the current (and then de-adapted) strategies.

This also suggests a more complex relationship between resilience and anticipation. It links resilience not only to the anticipation of what may happen, but also to the anticipation of coping capacities. This implies the monitoring of the *current* degree of control of the situation, and the prediction of the *future* level of control. As Woods puts it in Chapter 9 'To be resilient, a system always keeps an eye on whether its adaptive capacity, as it currently is configured and performs, is adequate to meet the demands it will or could encounter in the future.' And because the adaptive capacity includes anticipation, we have a recursive relation, namely that resilience also implies anticipation of future anticipation capacities. This is illustrated by the US Airways Flight 1549 captain's decision not to attempt to return to an airport, because such a decision would have engaged an irreversible course of action, with a total loss of control and inescapable catastrophic consequences if wrongly taken. This idea of constantly monitoring the future *marges de manoeuvre* and adapting the current state of affairs to protect these margins is also present in the background of what Bergström and his colleagues (Chapter 4 this volume) address in their experiment for '*Training organisational resilience in escalating situations*.' They outline a theoretical framework of 'Generic Competencies in Management of Escalating Situations,' which includes *the ability to constantly monitor whether the organisation is suited to manage the situation at hand*, and the ability to constantly *monitor and update the process* by which the escalating situation is managed.

Being Prepared to Be Unprepared

However, the conditions for resilience cannot be reduced to anticipation. While things that happen are controllable only if they have been anticipated to some extent, they will never have been anticipated in every detail. Hence, resilience implies a combination of readiness and creativity, and of anticipation and serendipity. Or, to put it differently, a resilient system must be both prepared, and be prepared to be unprepared. There is an explicit double bind in this last sentence that goes beyond a play on words. In the Hudson case study, the author shows that anticipating strategies at the system level in aviation can lead to a paradoxical result of generating unprepared operators. Building on Mintzberg's (1996) idea of 'predetermination fallacy,' he suggests that there may be an 'irony of resilience' in the fact that the real time competences needed to cope with unanticipated or extreme events are exactly those that are lost in the continuous attempt to anticipate all events and to pre-determine corresponding responses. This raises at least two issues.

The first one is the relationship between training and 'fundamental surprise' situations (Lanir, 1986). At first glance, it seems paradoxical to want to train people for something unimagined. How can one possibly train for such situations? In '*Lessons from the Hudson*' (see also Pariès and Amalberti, 2000), the author argues that it is possible indeed. However, the efficiency conditions (Dekker et al., 2008), including emotional and cognitive fidelity simulation of 'real surprise' are not met by the current aviation training system. One reason may be the dominant aviation safety paradigm which assumes that flight operations can be entirely specified by procedures that consequently must be fully adhered to by crews (cf., Chapter 18). Bergström and his colleagues also believe it is possible (Chapter 4); they designed scenarios aiming at training generic team competencies, rather than domain-specific skills representing pre-defined responses. They experimentally show the potential benefits of such training programs in a class of trainees to become incident commanders in rescue services.

The second issue is the relationships between resilience features engineered into the system as a whole, and resilience features of the local components or agents of the system (e.g., front line operators). Resilience at the global system level can be seen as a property emerging from the interactions of individual agents' behaviors, while at the same time, resilience at the individual behavior level at least partially is an outcome of the global system design. So, 'real time resilience' is generated through both a bottom-up and a top-down process. Woods (2006) refers to 'downward and upward resilience' to describe these interrelated and complex processes. Here we will only discuss the risk that anticipating strategies at the system level may generate unprepared operators. One way to help overcome this apparent double-bind is to adapt the level of functional abstraction of the prepared responses to match the level of uncertainty associated with the anticipated situation. In the Hudson River ditching case, every time the next line of the 'defence in depth' is breached, the uncertainty increases: the situation gets more improbable and less controllable. At each stage, the prepared responses shift from concrete and detailed to abstract and generic, the procedures shift from accurate and detailed action oriented protocols to a generic, goal-oriented response framework.

Adapted or Adaptive?

However, it would be naïve to pretend that merely adapting the level of functional abstraction of the prepared responses can solve the potential contradiction between anticipation and serendipity. In his theory of cognitive adaptation, Piaget (1967) argued that it included the two processes of assimilation and accommodation. Assimilation 'filters' the world to make it fit to individual mental structures. Accommodation works in the opposite direction: it modifies the filter (one's mental structures) to fit to the demands of the environment. Assimilation is mainly supported by homeostatic routines, while accommodation requires self-modification of existing schemes triggered, for instance, by cognitive dissonance. Similarly, the adaptive capacity of a system derives from its permanent under-adaptation (dissonance), which creates the tension for adaptation. Indeed, a complex system is necessarily partially 'out of tune' regarding its environment. Implementing its adaptation capacities to fit internal and external changes increases its spectrum of potential behavior, which momentarily provides the solutions to match the new needs (adaptation), while it increases the adaptation repertoire, and generates new exploration capacities, which will eventually expose the system to face new situations, and new challenges (de-adaptation). Finally, all this boils down to the optimality/brittleness trade-off constraining the behavior of any complex adaptive system (Doyle, 2000). Such a system cannot be at the same time totally adapted to its environment (optimally performing) and able to cope with disruptive changes in that environment (resilient). Adapted or adaptive, an inescapable choice has to be made.

Chapter 2
Lessons from the Hudson

Jean Pariès

> Things that have never happened before happen all the time.
>
> Scott D. Sagan (*The Limits of Safety*)

The successful ditching of US Airways Flight 1549 into the Hudson River (15th January 2009) shows an implementation of the 'strategic resilience' engineered into the aviation system – in this case, the multiple layers of 'defence in depth' set up by to manage a total engine failure in the context of bird ingestion. Each move through one line of defence to the next is like a tactical retreat, in which sights are lowered and sacrificing decisions are made, in order to save what can be saved. At each stage, the situation gets more improbable, more variable and less controllable; the probability and potential magnitude of damage are increasing; the response options are more restricted, harder to anticipate, more constrained by time, and less reversible. So the 'tactical retreat' is also a shift from adaptation to de-adaptation. Hence resilience implies a combination of anticipation and serendipity: a resilient system must be *both prepared and prepared to be unprepared*. But there may be a negative interference between anticipation and serendipity, leading to an '*irony of resilience*': the 'real time' competences needed to cope with unanticipated or extreme events at the 'sharp end' are exactly those which are lost in the continuous attempt to anticipate all events and to pre-determine corresponding responses at the system level. So there is a trade-off between efficiency (linked to the degree of adaptation of a system) and flexibility (linked to the adaptation bandwidth of a system).

Miracle on the Hudson River?

Many readers will probably remember the breathtaking images of the US Airways Flight 1549 ditched into the Hudson River, in New York, on 15th January 2009, with the passengers standing on the wings of the floating airliner. The entire crew was awarded, among other honours, the Master's Medal of the Guild of Air Pilots and Air Navigators (GAPAN). The GAPAN citation read: '*This emergency ditching and evacuation, with the loss of no lives, is a heroic and unique aviation achievement*.' Statistically, the event was indeed rare. There have been only very few documented occurrences of controlled ditching by commercial public transport aircraft. And it appears that prior to our recent US Flight 1549, only one known ditching of a passenger jet had been managed without fatalities (in St Petersburg, Russia, in 1963, an Aeroflot Tu124 jet ran out of fuel during an emergency and landed on the Neva River. All 52 people aboard survived and the jet was towed to shore).

So was the Hudson River successful ditching a miracle, a heroic achievement, or more simply an expression of some response capabilities fundamentally engineered into the aviation system, namely its *resilience*, as described by Hollnagel et al. (2006). And if the latter assumption is true, can we fish some good lessons from the Hudson River concerning resilience?

The Bird Hazard

The main hazard at the heart of the Flight 1549 accident scenario is what aviation people call 'bird hazard': in-flight collisions with birds, damaging the aircraft engines or the aircraft airframe. Bird strikes are exactly the opposite of unexpected events: they have been a well recognised aviation problem for decades. Actually, the first reported collision between an airplane and a bird is apparently as old as aviation: it happened to Orville Wright in 1905 (Bird Strike Committee-USA, 2009)! Things got considerably worse with the introduction of jet aircraft. Higher speeds increased the impact energy and jet engines demonstrated both a strong tendency to inhale birds and a chronic fragility to their impact. With the growth of aviation, bird strikes have become a very common event. From 1990–2004, over 56,000 bird strikes to civil aircraft were reported to the Federal Aviation Administration (FAA) in the US. This is considered to be a mere 20 per cent of the number that likely occurred. Worldwide, the bird strike damage cost to civil aviation is estimated to be over one billion dollars. And the issue is not only losses in dollars. In the 1990s the US Bird Strike Committee estimated that there is a 25 per cent chance in any decade that birds could cause a major airline crash.

In summary, at the aviation system level, bird strikes are a well known, frequent event and the associated threat for flight safety is well understood and recognised. So why would one refer to 'resilience' for such a well anticipated disturbance? The quick answer is that birds cannot be kept away from the daily environment of aircraft flights, while at the same time, as the Hudson ditching illustrated, birds can inflict very severe damage to aircraft. So what is well known, frequent and anticipated at a global system level can still be a big and challenging surprise at the front line operation level.

Bird Strike Protection Strategy

A lot of effort has been made during the last decades to reduce the frequency of bird collisions by controlling wildlife where possible, in airports and their vicinity. There are a number of techniques that can reduce the number of birds in the vicinity of airports: making the environment unattractive for birds, scaring the birds or reducing the bird population. None of them can really solve the problem. An airport is a part of the local ecosystem and eliminating any one problem species will only lead to some other species taking its place. Furthermore, while the majority of encounters with birds occur at low altitudes (below 1500 feet) in the vicinity of airports, they also happen at higher altitudes and even – but very rarely – at cruise level (the world height record for a strike is 37,000 feet!). Consequently, birds are – and must be considered as – a usual component of flight environments, hence bird strikes must be considered as 'normal' and frequent events.

Unfortunately, aircraft are not really immune to birds. Trying to make aircraft more resistant to bird strikes has been a permanent target of airframes, wind-shields and engines certification standards for decades. Large commercial aircraft are certified to be able to continue flying after impacting birds. Jet engines are designed to withstand bird strikes, at least to some extent. They must demonstrate their ability to cope during a series of certification tests during which two-kilogram chickens, a series of eight 'medium' size birds and sixteen 'small' birds, are shot out of a cannon at their blades while running at full power. So engine blades are extremely tough, and aircraft engines routinely ingest birds without any damage. But these tests have serious limitations. Some Canadian geese subspecies weigh over 7 kg. And most of these large birds travel in flocks. No engine could be made both acceptably efficient and immune to a flock of Canadian geese. There is a trade-off here between safety and fuel-efficient air travel. As Downer (2009) puts it, 'A pant-wetting splash in the Hudson once a decade is probably an acceptable trade for cheap and fuel-efficient air travel.' And even flocks of small birds (e.g., starlings) can cause engine failure and substantial damage, when they are present in large numbers.

In other words, a jet engine cannot be designed and certificated to resist even rather probable encounters, not to mention the worst case scenario of flocks of large birds. So the next line of defence is set up, at the crew-aircraft system level, as follows: if any one engine is unable to continue generating thrust, the airplane will get enough power from the remaining engine or engines to safely reach an accessible runway. Hence, a commercial aircraft equipped with N engines can only be certificated if it keeps acceptable flight performance – including climb capability – with $(N - 1)$ engines.

But when the Flight 1549 twin engine Airbus A320 carved into a flock of Canadian geese about two minutes after take-off, at about 3000ft, *several* of these huge birds were sucked into both engines. This clearly exceeded the above mentioned engine certification criteria, as well as the $(N - 1)$ engine performance condition: both engines suffered a simultaneous and sudden loss of thrust. Is there then a next line of defence available? Can it be made reasonably certain that a dual engine failure on a twin engine aircraft will only result in a 'pant-wetting splash in the Hudson'? The answer is not straightforward.

On the one hand, a total loss of thrust is anticipated in the aircraft certification principles. Emergency flight-management resources are provided. Several systems and procedures are made available to ensure that the crew can continue to maintain some aircraft control (even if, in the case of vertical speed, this control is obviously limited). The APU (Auxiliary Power Unit – a small turbofan engine fitted in the tail of the aircraft) can be used (and was actually used by the US Airways crew), as a spare electrical power supply. A RAT (Ram Air Turbine) can be deployed (and was deployed) to produce the ultimate hydraulic pressure needed by the flight controls. Emergency procedures are provided to properly manage the remaining resources. As a matter of fact, on board flight 1549, the normal electrical supply and all three hydraulic systems remained fully operational, and the flight control law remained in 'normal law' at all times until the ditching. (The 'normal law' is the flight mode available when the whole flight control system is functioning normally: the side-stick deflection then controls the load factor independently of speed and aircraft configuration, pitch trim is automated and flight envelope protection is provided throughout the envelope.)

On the other hand, all of these emergency flight management resources can only be aiming at one thing: keeping enough control on the aircraft to pilot an inescapable descent. This is of course critical: a loss of control in flight would inevitably lead to tragedy. But what will happen at the end of the 'controlled' descent is much more difficult to control and most often falls into Westrum's 'category 3' situation (Westrum, 2006). It heavily depends on the circumstances, that is, when and where the dual failure occurred. The situation will be very different if the engines quit during daylight, in good visibility conditions, at an altitude and distance such that the aircraft can glide to an airport and end up on a runway, than if they fail at low altitude in the middle of nowhere, amid mountains, above the sea, or above a large city. So the final outcome will depend heavily on 'providence,' as well as on the pilots' gliding skills and ability to make the right decisions.

Additionally, the odds for a thrust-loss scenario over water are far from negligible, hence ditching has also been anticipated.

There is a ditching procedure included to guide the crew actions. There is a set of aircraft design features intended to limit damages to the hull and facilitate the aircraft's ability to float in case of ditching, including a 'ditch button' closing all valves to make the cabin watertight. There are cabin procedures for ditching and evacuation (including the routine life jacket briefing that most of us pay little attention to while settling back into our seat).

But landing a large jet on water remains a highly unusual and hazardous operation. In the case of Flight 1549, the ditching occurred 3 minutes and 30 seconds after the thrust loss. The ditching airspeed was about 130 knots – just a few knots above the minimum speed with flaps and slats in configuration 2 and landing gear up. Pitch attitude was 10 degrees nose up and the wings were perfectly level, a critical condition to avoid a potentially devastating asymmetrical impact. The aircraft fly-by-wire design and its embedded stability and stall protection and quite a large dose of luck also contributed. It was daylight, there was a clear sky and good visibility, a river rather than the open sea nearby, a smooth water surface with only a light surface wind and the crew were familiar with the area. The aircraft ability to float was reduced by severe damages to the tail generated by the impact with water with a high nose up attitude, but it floated long enough for all the occupants to be safely evacuated. The evacuation of the aircraft was a nice piece of effective cooperation among and between the cabin and cockpit crew. No boats were hit during the landing, but many were readily available at the scene to assist with the rescue.

From Anticipated Emergency to Real Time Response

What has been described above may give a feeling that the management of such an event is totally based on anticipation. As a matter of fact, even when they are well anticipated at the overall system level, emergency situations like this always come as a 'fundamental surprise' (Lanir, 1986) when they jump on front line operators in a handful of seconds. They immediately trigger a major cognitive conflict between current mental representations and current experience, first leading to shock and denial.

- *'It was the worst sickening pit of your stomach, falling through the floor feeling I've ever felt in my life. I knew immediately it was very bad'* Chesley Sullenberger, the US Airways Flight 1549 Captain, told CBS News. *'My initial reaction was one of disbelief. I can't believe this is happening. This doesn't happen to me.'*
- Patrick Harten, the New York TRACON La Guardia Departure Air Traffic Controller who last spoke to Flight 1549, gave a testimony of his reaction to a Congress Panel (Committee on Transportation Infrastructure, 2009): *'I asked him to repeat himself, even though I heard him just fine. I simply could not wrap my mind around those words.'* And when the plane disappeared from his radar screen: *'It was the lowest low I had ever felt; Truth was, I felt like I'd been hit by a bus'* he said.

So a key question is: what capacities are needed of front line operators to properly respond to such situations? Are they specific to the emergency, or are they simply the extended implementation of daily adjustment abilities? We can see a whole set of skills and response capacities behind the Hudson success story. In spite of the initial denial phase, a very fast overall 'operational' comprehension of the situation was achieved. The assessment of the options was globally right and this was instrumental for the success: any decision in such a situation is a single shot. There is no 'please try again' button. It took a subtle balance between experience (e.g., building on past glider experience) *and* opportunism (e.g., taking advantage of the fortuitous presence of the Hudson amid a highly populated area; choosing a ditching location near operating boats so as to maximise the chances of rescue), self confidence (*'I was sure I could do it'* said Captain Sullenberger in a post-event interview) *and* awareness of limitations (a right mixture of 'yes we can' and 'unable'). It took a highly dynamic (re)planning capacity, allowing for good-decisions-in-series. According to the transcript of the communication recorded between the crew and the La Guardia Departure Controller, the first intention declared by Captain Sullenberger when the engines failed was to return to La Guardia. He was then offered runway 13 by the Controller but realised it would not be possible to return to La Guardia. He then briefly considered going to Teterboro Airport and rejected this option as well.

The successful ditching also required the following of (emergency) procedures, as well as interpretation, adaptation, improvisation (e.g., the First Officer started the APU, although this was not required by the procedure) and even some kind of 'bricolage': a relight attempt procedure started, not finished due to lack of time, the ditching procedure started, not finished by lack of time, and so on. It also took quick and efficient communication among the crew, as well as between the crew (cockpit and cabin), and between the crew and the Air Traffic Controller.

A key issue was controlling stress. In the previously mentioned interview, Captain Sullenberger said: '*I was not this calm then, but I was very focused.*' In his testimony the Air Traffic Controller used the same words to describe his state of mind during the event: '*During the emergency itself, I was hyper focused, I had no choice but to think and act quickly, and remain calm.*' And '*I was flexible and responsive, I listened to what the pilots said, and made sure to give him the tools he needed. I stayed calm and in control.*' Actually, training and experience are the key issue in this perspective: when people are well trained and experienced, they become highly focused and do not fall apart. The US Airways crew was very experienced. Captain Sullenberger had more than 19.000 flight hours, his First Officer more than 15.000, the three Cabin crew had between 26 and 38 years of experience.

Incidentally, this event also illustrates a perspective on resilience that is not directly related to systemic resilience, but is nevertheless worth mentioning here: individual (psychological) resilience to such a traumatic event. The hardest part is the post-event period. As the Traffic Controller put it: '*It may sound strange, but to me the hardest, most traumatic part of the entire event was when it was over … when it was over, it hit me hard.*' And also: '*Even when I learned the truth, I could not escape the image of tragedy in my mind. Every time I saw the survivors on television, I imagined grieving widows. It's taken over a month for me to see that I did a good job.*' He could only return to his job after 45 days of paid leave, admitting that '*it may take time for me to regain my old confidence*'. Ironically, the Controller seems to have been more severely affected than, for example the cockpit crew. As a matter of fact, this is perfectly predicted by models of psychological resilience, which emphasise the role of three main enablers: involvement into action, altruism and group solidarity.

From 'Satisficing' to 'Sacrificing' Decisions

There have been some discussions after the event about the fact that the US Airways crew did not activate the 'ditch button' to make the cabin watertight. Beyond the fact that it would not have made a big difference, due to the damages already suffered by the airframe, it is very important to understand that, in such crisis situations, there is no error-free course of action.

In his interview by US TV presenter Larry King, Captain Sullenberger stated: '*I expected that this was not going to be like every other flight I'd flown, for my entire career and it probably would not end on a runway with the airplane undamaged.*' This statement shows the huge amount of uncertainty that suddenly prevails in such moments, both about the current state of affairs and its short-term evolution. Suddenly, the known and anticipated world was lost. The actual status of the situation was blurred. The future was suddenly heavily uncertain, and at the same time, critically depending on (irreversible) decisions, which had to be made very quickly, and could only be based on fast judgements and generic risk-assessment with very few factual and procedural references. The lack of information about the actual situation and its potential evolution is such that decisions are never completely 'right' when analysed with the benefit of hindsight. Everyday human behaviour has been described (Simon, 1982) as 'bounded rationality'. In crises situation, the bounds get far worse. The lack of time and knowledge and resources becomes overwhelming. 'Satisficing' (rather than maximising) decisions become 'sacrificing' decisions. The scenario of US Airways Flight 1549 ditching provides several illustrations of this.

A first example was the decision to ditch in the Hudson River itself. Captain Sullenberger said in the above mentioned interview: '*I quickly determined that we were at too low an altitude, at too slow a speed, and therefore we didn't have enough energy to return to La Guardia, because it's too far away and we headed away from it. After briefly considering the only other nearby airport which was Teterboro in New Jersey, I realized it's too far away.*' Post-event simulations showed that considering its speed, distance and altitude, the aircraft had actually enough energy to glide back to the La Guardia runways. But this is what we know after hours of data processing and simulations. There was nothing available to the crew to determine accurately and reliably whether or not they could make a runway. Their only reference was their experience, their airmanship, their feeling of the situation. '*… And the penalty for choosing wrongly, and attempting to make a runway I could not make might be catastrophic for all of us on the airplane plus people on the ground.*' So there was an implacable trade-off here: either the Hudson, certainly bad but possibly not catastrophic, or surrounding airports, possibly happy ending, with minimum damage to the airplane, but almost certainly catastrophic in case of failure of the attempt. What we have here is a nice piece of risk management, through a 'sacrificing decision': minimising the odds of a disaster by deliberately sacrificing the most ambitious, potentially happy ending – but intolerant – branch of the options tree, to set up a kind of bottom line class of damage, associated with a ditching.

A second example was the engine relight attempt made by the crew. Both engines suffered a simultaneous and sudden loss of thrust. The A320 'Dual Engine Failure' emergency procedure calls for relight attempts. At about 500ft and 200kts, the crew attempted a quick relight on engine 1, without success. But while there was no further response from engine 2, engine 1 continued to deliver some thrust (about flight idle thrust) for about 2 minutes and 20 seconds. It was still slightly moderating the rate of descent and producing hydraulic and electrical supply. A damaged engine can in some cases restart in a better thermodynamic regime (a bit like a jammed computer after a reset). But relighting also implies to stop the engine first, with the risk that it will not relight, with possibly a regression of the flight control mode to 'direct law'. Unlike the 'normal' or 'alternate' laws, the 'direct law' is an emergency flight mode, available as a back up when the flight control system is heavily damaged. In this mode, artificial stability and flight envelope protection are lost and the control of the aircraft would have been much more difficult, particularly during the flare. So there was again a risk trade-off here: risking partially losing the available controllability of the aircraft, in order to potentially regain enough thrust to reach a runway, as every airline pilot knows that landing a large jet outside a runway is a potentially catastrophic event. Interestingly, what we have in this case is the opposite of the previous 'sacrificing decision'. Once the bottom line damage was insured by selecting the ditching option, attempts were made to re-activate the 'make-a-runway' option (without its terrible penalty in case of failure), by attempting to relight at least one engine.

From Safety Strategies to Resilience Engineering at the System Level

While it does not officially uses the word 'resilience' (yet), the aviation community has its own way of recognising a hierarchy of disturbance, from normal to totally abnormal situations. Three main 'operational safety domains' are commonly referred to:

1. A '*normal operation*' envelope inside which, more or less, the course of events follows pre-defined tracks, people follow procedures or expected behaviour, parameters stay within design limitations, and variations are compensated and absorbed by intrinsic flexibility within the system. In this normal flight envelope, real birds are not worse than the certification test chicken, engines can sustain their impact, and flights make it to destination on time (or nearly). The crew-aircraft system is 'adapted' to the situations encountered, which means that the pace and magnitude of variations and disturbances stay within its routine adjustment or tolerance capabilities.

2. An '*abnormal operation*' envelope inside which the course of events significantly departs from normal tracks, although mainly in anticipated ways: parameters exceed design limitations, important components (e.g., engines) fail, and the like. These disturbances need to be actively handled. They are managed by specific procedural responses, intentionally built-in redundancies, intrinsic resistance, or flexibility within the system. In this 'abnormal flight envelope,' real birds are worse than in the certification test, one engine may suffer from their impact, and the flight may return back to the departure airport or have to divert. The pace and magnitude of variations and disturbances are such that the crew-aircraft system needs to quickly 're-adapt' itself to the situation, through a pre-defined and active reconfiguration process.
3. Finally, an '*emergency operation*' (open) region inside which the course of events departs from normal tracks in extreme proportions and possibly totally unanticipated ways: parameters 'go crazy,' critical components are lost, and exceptional disturbances threaten the possibility of keeping control on the flight. These emergency situations urgently need to be managed, by specific or generic procedural responses, creativity, built-in ultimate backups, or intrinsic toughness. In this 'emergency flight envelope,' real birds are far worse than in the certification test, and can shut off all engines. The flight may well end up prematurely far from its expected destination. The crew-aircraft system is both unable to 're-adapt' itself to the situation, and able to keep some form of control on it, and mitigate its consequences, provided it can quickly and fully stretch its relevant capabilities.

While the notion is not commonly used in aviation safety language, one could recognise behind this hierarchy of responses a compact version of a classical 'defence-in-depth' strategy. For example, concerning bird hazard, the first line of defence is minimising the frequency of bird strikes. The second is the ability of the aircraft and its engines to hit some birds without damage.

The third is the ability of the crew-aircraft system to continue flying and to reach a (possibly alternate) airport after impacting birds and losing one engine. The fourth is the ability of the crew-aircraft system to keep enough flight controllability after losing all engines to be able to glide towards a runway or any suitable crash-landing area. The last line of defence is the ability of the crew-aircraft system to land on unprepared terrain or ditch with minimum damage and to safely evacuate passengers.

While not a new concept, such a defence-in-depth structure can be regarded as a '*strategic resilience*' engineered into the system at the highest, holistic level. Each move through one line of defence to the next one means a kind of tactical retreat, in which sights are lowered, sacrificing decisions are made, in order to save what can be saved from the wreckage. At each stage:

- the situation gets more improbable, more variable, less controllable;
- the probability of damage is increasing, as well as the potential magnitude of damage;
- response options are restricted, harder to anticipate, less reversible, deadlines are more demanding.

So this *series of tactical retreats* is also a march from well known territories to unknown areas, *a shift from adaptation to de-adaptation*. Associated operating procedures shift from accurate and detailed action protocols to generic frameworks. Front line operators need to refer to different models of the world, from which they need to derive both an overall sense-making (Weick, 1993) and 'satisficing,' then sacrificing solutions. Built-in extra robustness must also be provided in order to allow the system to cope with anticipated or unanticipated stress (e.g., sustaining a landing on water for the fuselage).

An important challenge is then to better 'engineer' these capacities into the system. At the whole system level, this could be obtained by doing more of what is already done: for example, increasing the efforts of controlling wildlife on airports and their vicinity, developing ground-based and airborne bird detection and avoidance systems, continuing to harden engines blades, and so on. But the main change would be qualitative: addressing the emergency situation *management* needs as such, at a systemic level (aircraft design, operational procedures, crew training, Air Traffic Control procedures, etc.). For example, it would not even take a second for the Flight Management System to answer a question like 'can we make Teterboro?' (and many others) much more accurately than crew intuition, if a specific module (an all-engine-out emergency management assistant) was incorporated.

When Systemic Resilience Efforts Undermine Resilience at the Sharp End

But the best strategy is worth nothing if it cannot be implemented at the 'sharp end' of the system, by front line operators equipped with the corresponding skills. So, is the current airline pilot training system efficiently providing the competences needed for a resilient system? The worldwide celebration of the heroic behaviour of Flight 1549 crew speaks by itself: Captain Sullenberger's skills have been attributed to his personal talent, his past experience as a glider pilot, rather than to his specific airline pilot training.

The 'ACCOMPLI' project is a three year (2006–2009) research project funded by the French DGAC, aimed at developing pilots training approval criteria based on pilot cognitive competence development. A literature review of competence models was compared to what industry professionals (airline managers, pilot instructors, and airline pilots) describe as their perception of co-pilot competence needs. *Ab initio* training was observed from entry to line adaptation, with the aim of spotting competence building processes. Within that project, young First Officers from several European airlines were asked to fill a web-based questionnaire to identify which aspects of their job they felt confident with. All their answers mentioned flying the plane in normal situations (while they recognised at least a six-month adaptation as necessary for confidence), and handling anticipated and trained 'abnormal' situations. When asked what they did *not* feel confident with, all answers mentioned relationships in the cockpit (handling the diversity of – real – Captain personalities and practices towards company procedures) and most of them mentioned managing:

- the diversity of operational situations, with real surprises;
- 'the edges of Standar Operating Procedures (SPO),' induced stress when over the limits of SOPs;
- borderline situations (e.g., non stabilised approaches in a dense traffic);
- interpretation and adjustments of procedures to current situations;
- interruptions, multi-tasking and interactions;
- the diversity of accents and foreign phraseology in ATC.

Obviously, what First Officers discover in their first months of airline operations is that the real world is much more varied, unstable, uncertain, surprising and complex than the one with which they are confronted during training. When asked for improvement suggestions, they wish they had been given more generic tools to help face the unexpected and managing surprises. They suggest more training in the simulator for operations that are not 'clear-cut' but rather at the edge of normal operations. They want better feedback on actual operations, incidents and accidents worldwide, with realistic 'emotional examples.' They suggest cultivating 'common sense' and 'airmanship' by being confronted with the unexpected during training by experienced instructors ('*teachers make the difference, and not the programme*'). They suggest flight simulator sessions where pilots would *not* been briefed before, and would not know in advance what they are about to be trained for.

These suggestions by young First Officers were easier to understand within the ACCOMPLI project from the perspective of expected cognitive competences. The following list of such competences was established through a review of literature, and questionnaires and interviews of pilots, flight instructors, training organisation managers, and airlines managers.

- To be able to construct and maintain an adequate (distributed) mental representation of the situation.

- To be able to assess risk and threats as relevant for the flight.
- To be able to assess self proficiency envelope, to know and recognise its boundaries, and to adapt one's tactics and strategy accordingly.
- To be able to switch from a situation under control, to a crisis situation (recognition, coping).
- To be able to construct and maintain a relevant level of confidence (towards self, others, technology).
- To be able to learn, implement and maintain the routines and skills associated with the basic flight functions (fly, navigate, communicate).
- To be able to contribute to / to make a decision in a complex (uncertain) environment.
- To be able to manage interactions with aircraft automated systems.
- To know, to understand and to be able to speak the aviation 'jargon'.
- To be able to manage interactions with, and cooperate with, crew members and other staff.
- To be able to make an intelligent usage of procedures.
- To be able to use available technical and human resources, and to re-configure them.
- To be able to manage time and time pressure.
- To be able to properly transfer acquired knowledge and know-how from a specific context to a different one.
- To be able to properly use and manage information and communication technology equipment (ICT).

These dimensions of pilot (cognitive) competence are neither mutually independent, nor do they present a specific order or hierarchical relationship. They can therefore not really be developed or implemented in isolation from each other. However, it can be easily guessed when looking at this list that a significant part of them – namely the first five, which are particularly important for unexpected or abnormal situations – will not be found in the explicit training objectives of most pilot training syllabi. They are a *side product* of the training rather that its central target. Even when 'abnormal' operations are addressed as such, the training considerably lacks realism. The syllabi are currently so time-constrained that each exercise, including emergencies, is pre-briefed, practised only once without the surprise factor, and debriefed, with little or no opportunity to reinforce lessons. The role of solo flight training for multi-crew pilots is more and more contested, although it was perceived by the interviewed First Officers as essential for building confidence and decision-making skills, as well as self-knowledge and personal limitations awareness, because it confronts trainees to the unexpected, and to the emotional context of vital decisions.

In other words, there seems to be a weakness of the current training system in building proper uncertainty management competences, and this weakness is far more than a side effect of the simplification of the world needed for any training environment. It is a direct result of a general training strategy that is just caricaturing a bit the very normative, procedural aviation safety strategy. Training efforts are completely oriented towards this very objective: learning how to recognise a set of anticipated situations, and how to respond properly to them with the relevant pre-established procedure. This leads to what Mintzberg (1994) calls the 'fallacy of predetermination': there is so much emphasis on anticipation and planning that there is no consideration any more for events that fall outside the anticipated envelope, and the illusion develops that the world will unfold as anticipated.

The bad news is that uncertainty management competences are not generated and maintained better by the handling of daily variations, which rather resorts to homeostatic adaptation routines. They develop from a recurrent confrontation with challenging, surprising, unexpected, threatening situations. As these are hopefully not encountered in daily operations (and the safer the system, the less they will be), substitutes must be found. An efficient feedback from incidental and accidental experience, as well as a proper use of simulation, can obviously provide such exposure. Within the European Commission Research Project 'ESSAI,' a specific one-day training has been developed for improving 'situation awareness' and crisis management capacities in the cockpit. Positive effects of the training could be demonstrated, particularly on situation awareness, and threat management. Mental simulation and counter-factual thinking have also been shown to be efficient. In a study on behavioural markers of 'procedural excellence' amongst a group of UK paediatric cardiac surgeons, Carthey et al. (2003) found that the best scores in terms of fatality rates or near misses were those of surgeons regularly practising 'what if' mental simulations of complications.

In Conclusion: Two Lessons and a Wish

The first lesson to be taken from the Hudson River ditching is that a resilient system must be *both prepared, and prepared to be unprepared*. Because things will happen that only are controllable if they have been anticipated to some extent and, at the same time, that will never be anticipated in detail. It means that we need a) generic anticipation schemes, providing (common) sense-making frameworks of what happens, at a level of abstraction which is high enough to wrap around all the countless and unpredictable variations of real stories, and b) fast and efficient implementation sketches and skills, capable of forcing the available generic schemes to fit the parameters of the day, under critical time constraints. In other words, resilience implies a combination of anticipation and serendipity. An efficient way to engineer this combination into a system is to set up a hierarchical defence-in-depth strategy, where a breach in a line of defence triggers a tactical retreat behind the next one, with operating procedures shifting from detailed protocols for normal situations, to a generic action framework for emergency situations.

A second lesson is that there may well be something wrong in the relationship between 'downward' and 'upward' resilience (Woods, 2006a), something like an 'irony of resilience,' similar to the '*irony of automation*' described by Bainbridge (1987): the competences suddenly needed at the sharp end to cope with unanticipated or extreme events are exactly those which are lost in the continuous attempt made at the blunt end to anticipate all events and pre-determine corresponding responses, or eradicate the extreme. Proceduralisation and automation both try to reduce the uncertainty in the system by reducing variety, diversity, deviation, instability. But the side effect is that this also reduces autonomy, creativity, and reactivity. Increasing order, conformity, stability, predictability, discipline, anticipation, makes the systems better (more efficient, more reliable), possibly cheaper and generally safer within the confines of their standard environment. They also make them increasingly brittle (less resilient) outside the boundaries of the normal envelope. We have to recognize that there is a universal trade-off between efficiency (adaptation degree) and flexibility (adaptation bandwidth). Desert lizards are so well adapted that they can survive for years without water, but would disappear if the climate changed by a few degrees.

So, until Resilience Engineering has found a way to overcome the irony, we can only hope that flocks of Canadian geese chose clear skies and very senior pilots when they fly across airport departure paths.

This page has been left blank intentionally

Chapter 3
Coping with Uncertainty. Resilient Decisions in Anaesthesia

Lucie Cuvelier and Pierre Falzon

This study aims to describe the variability anaesthesiologists deal with in their everyday work and to understand the different strategies used by them to avoid the negative consequences of this variability. An empirical research, based on the critical-incident technique was conducted in a paediatric anaesthesiology service in a French hospital. The results highlight a distinction between potential situations in which the problem was envisioned beforehand by practitioners and unthought-of situations which were unthinkable for the anaesthesiologist previously and at the time of their occurrence. This subjective classification based on 'the astonishment of the perceiver' highlights two critical decisions made by anaesthesiologists in order to manage variability. The first is the preoperative definition of an envelope of potential variability of the surgical intervention. The second concerns the occurrence of an event which trespasses the envelope initially defined and requires the mobilization of additional resources. The identification of these two critical decisions provides opportunities for researches and actions to enhance resilience in the practice of anaesthesia.

States of Resilience and Uncertain Events

Resilience is the intrinsic ability of a system to adjust its functioning so that it can sustain required operations under both expected and unexpected conditions. This latter definition proposed by Hollnagel (Prologue of this book) emphasises the breadth of the concept, indicating that resilience is not only the system's ability to cope with unforeseen variability that fall outside the expected areas of adaptations (Woods, 2006a) but also looks at its ability to operate in foreseen conditions. Indeed, since the appearance of the term resilience in the field of safety, many models have attempted to characterise the systems' domains of variability, in order to clarify this ability (or these abilities) to be resilient. These models generally cut 'the space of possibilities' according to the frequency of the disturbances that a system or an organisation may be facing. Indeed, a system that is resilient when a regular threat occurred will not necessarily be resilient facing an irregular threat or an unexampled event (Westrum, 2006). Thus, states of resilience can be defined as those proposed by Hollnagel and Sundström (2006): the state of normal functioning, the state of reduced irregular functioning and the state of disturbed functioning. A resilient system is then one capable of detecting that the conditions have changed, to assure transition to another state and to operate in the new state of resilience achieved.

Other studies have modelled how systems make adjustments to address these different types of uncertain disturbances, notably by describing how these systems behave when they must operate under high pressure. By analogy to the physical model of materials 'stress and strain', Woods and Wreathall (2008) distinguish two adaptation areas.

1. The uniform or elastic region, in which the organisational response is proportional to the increasing stress. In this region, there are plans, procedures, training and resources provided to allow the system to adjust to demand. This area corresponds to the envelope for which the system is designed (Woods, 2006a).

2. The non-uniform or extra region, for which the organisation's responses can no longer proportionally cope with the increasing load: failures appear and performance deteriorates (Wears et al., 2008). In this region, additional resources are mobilised and local strategies are deployed by individuals and teams to deal with disturbances.

Similarly, based on the model of socio-technical systems proposed by Rasmussen (1997), Miller and Xiao (2007) describe two adaptation areas: the compensation area (that corresponds to the marginal zone of the Rasmussen's model) and the decompensation area. The marginal zone represents 'the system's ability to cope' (Rasmussen, 1997). In this area, 'opportunistic processes' outweigh the disruptions: a range of behaviours and resources are mobilised to maintain the operating system to a level of risk as small as possible. Such compensations can sometimes mask the presence and the development of dysfunction (Woods and Cook, 2006). Once the compensation mechanisms are exhausted, the system decompensates: parameters suddenly collapse and potential for failures increases. This decompensation may be chronic or acute (Miller and Xiao, 2007), or both at once (Wears et al., 2006). Therefore, the study of resilience requires describing these different classes of adaptive processes that allow a system to adjust its functioning so that it can pursue operation under varying conditions (Woods and Cook, 2006).

Describing How Anaesthesiologists Manage Uncertainty

In the context of research on patient safety in paediatric anaesthesia, we sought to describe the variability that anaesthesiologists deal with in their everyday work and to understand the different strategies used by them to avoid the negative consequences of this variability. Indeed, with a risk of a fatal accident less than 1 per 100.000, anaesthesia is now faced with the 'paradoxes of safe systems': to continue to progress on safety, it is necessary to change the nature of the system and to consider different safety measures (Amalberti, 2001; Amalberti et al., 2005). One way could be to envisage specific training methods, similar to those used in the aviation domain with simulators (cf., Chapter 8). The objective of this study is then dual: first, it seeks to identify the different types of disturbances that anaesthesiologists have to manage in their real work activity and, second, to highlight the resilience factors, that is, the strategies developed in practice by anaesthesiologists to allow the system to function despite these disturbances.

The study was conducted in a paediatric anaesthesiology service within a university hospital in France. After several weeks of open observations, we chose to deploy an *a posteriori* methodology: scenarios of real incidents were collected after their occurrence. Indeed, observations conducted have shown that (thankfully) few incidents occur in daily practice. And although resilience is an essential quality for any type of disturbance, it is recognised that '*these determining characteristics are often easier to note in the case of events of an unusual scale or severity*' (Hollnagel and Sundström, 2006). Indeed, the analysis of these incidents shows how the system behaves at the performance boundaries, that is, simultaneously how it adapts and adjusts to cope with disturbances and what are the limits of this adaptation (Woods, 2006a). The method chosen to collect *a posteriori* incidents was based on the critical incidents technique (Flanagan, 1954) and its main extension: the critical decision method (Klein and Armstrong, 2005; Stanton et al., 2005). Both techniques aim to raise 'salient episodes' in the practitioners' memory: during interviews, anaesthesiologists were asked to recall and describe incidents they had experienced or to which they participated. As defined by Flanagan an incident is critical if '*it contributes positively or negatively to the overall goal of the activity*' (Flanagan, 1954, p. 272). It is therefore theoretically possible to collect events that had a particularly beneficial impact on the success of the activity. But in practice it turns out that one obtains mostly negative events (Bisseret et al., 1999). Thus, to understand the system performance in general – both failures and successes (cf., Prologue) – our collection of scenarios focused exclusively on cases of near accidents where adaptations were successful. For many authors, it is preferable to collect recent events to get stories less distorted and more detailed (Bisseret et al., 1999; Ombredane and Faverge, 1955). But studies on episodic knowledge show that the age of the recalled events is totally unrelated to the vividness of memory. The vividness seems to be mainly connected with the emotional content of the memory (Bærentsen, 1996). Thus, in the method we deployed, no restriction was made *vis-à-vis* the age of the remembered episodes: we considered, as Flanagan suggests, that '*the incidents themselves contain evidence of the accuracy of the account. If complete and precise details are given, we can consider that the information is accurate*' (1954, p. 275).

Six trained anaesthesiologists were interviewed (four of whom had many years of experience). They were asked to recall near-accident situations, in which they were close to a severe problem for the patient, but where they managed to cope and get back to a stable condition. Each interview lasted for about one hour. The instructions given to participants for verbalisations were fully prepared, because it is a crucial point for the verbalisation processes: '*a small change in the instruction may affect the nature of the collected incidents*' (Flanagan, 1954, p. 277). The interviews were then semi structured: the interviewees spoke freely but were brought to address predefined themes. According to the critical decisions method, 'probes' previously selected, were used to obtain more information on cognitive processes and key decisions (Stanton et al., 2005). At the end of the interview, two open questions were asked in order to conclude.

> In your opinion, which factors enable a team of anaesthesiologists to cope with emergencies?
>
> Are some events easier to recover than others?

These questions aimed to make the respondents see the connections between the episodes mentioned during the interview and to compare them. According to the critical incident technique, data processing was mainly qualitative and subjective: it consisted in building a classification of events collected. This is the '*classification criteria and the values they take that made the outcome of the study*' (Bisseret et al., 1999, p. 127). Interviews were thus transcribed and a content analysis was performed.

Unforeseen Situations: Potential Variability and Unthought-of Variability

Twenty-two situations of near accidents, dating from ‘a few days’ ago’ to ‘20 years ago’, were recalled by anaesthesiologists during the interviews. They allow us, as a first step, to draw up some characteristics of near accidents marking the memory of anaesthesiologists. The first characteristic of these scenarios concerns the severity of the situation. All of the situations reported are situations where the patient’s life was at stake, for which the anaesthesiologists said they narrowly avoided the death of the patient. The second characteristic refers to the temporal dimension of scenarios remembered: most situations are acute situations where time passes very quickly, where monitored parameters are changing ‘brutally’ and where anaesthesiologists must act in urgency. A third feature concerns the emotional content of narrated episodes: in half the cases, interviewed practitioners spontaneously evoked memories of ‘fear’, ‘stress’, ‘concern’ or ‘anguish’. One last point relates to the unexpected nature of the situation: all recollected situations were characterised by practitioners as unexpected events. But a more detailed analysis of the scenarios shows that the concept of ‘unforeseen’ is vast and includes many different situations. It is also mentioned by some practitioners that ‘*there are levels of unpredictability*’ that some episodes are ‘*more or less predictable than others*’. Indeed, unexpectedness can arise in different ways. An unforeseen situation may be a situation that was already envisaged as possible by the anaesthesiologist before the intervention. In this case, the unexpected is not directly related to the event but to the time of the occurrence of this event, that could not be determined with certainty by the practitioner before surgery. These situations are *potential situations.* At the opposite, a situation may be unexpected in its very nature: the event itself has not been foreseen by the anaesthesiologists. The situation is not surprising because of its unexpected occurrence but because of its very nature, which has not been thought of. These situations were *unthought-of situations* when they occurred. The distinction between these two types of unforeseen situations defines two areas.

- The area of *potential* variability, which corresponds to the situations that the anaesthesiologist considers a priori as likely to occur during surgery.
- The area of *unthought-of* variability, which corresponds to the situations which are not envisaged by the anaesthesiologist before surgery.

Unlike the 'adaptation areas' and the 'states of resilience' presented above, this classification is subjective since it relies on a categorisation of situations as they were experienced by the subjects: it takes into account the operators' point of view on the events, in a given context, when the event occurred. This distinction is based on an ergonomic activity analysis, that is, on 'an analysis of the strategies (regulation, anticipation) developed by operators to manage the gap between the prescribed task and the actual work' (Guérin et al., 1997). The goal is not to characterise the rarity or the 'objective complexity' of the event as seen by a subject outside the action (Leplat, 1988) but to describe '*the astonishment of the perceiver*' (Weick, 1993, p. 633) and therefore, this classification takes into account the way in which the situation was envisaged before surgery. Figure 3.1 places these two areas in the model of 'the resonance in complex systems' proposed by Hollnagel (2004).

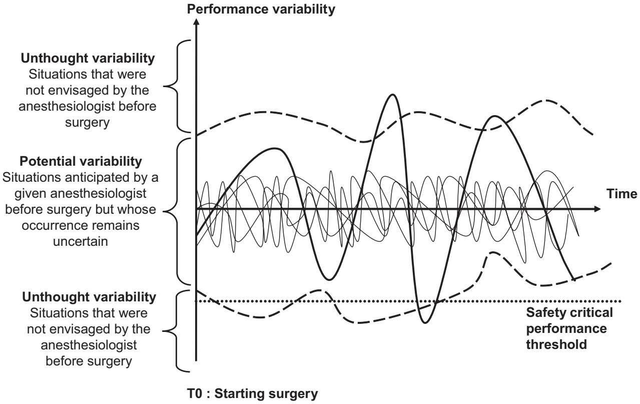
Performance variability
Unthought variability
Situations that were not envisaged by the anesthesiologist before surgery
Potential variability
Situations anticipated by a given anesthesiologist before surgery but whose occurrence remains uncertain
Unthought variability
Situations that were not envisaged by the anesthesiologist before surgery
Time
Safety critical performance threshold
T0 : Starting surgery

Figure 3.1 Schematic representation of the categorisation of collected episodes

Note: each episode is classified either as an unthought-of situation, or as a potential situation, according to the operator's perspective in the incident's circumstances

This distinction between *potential* and *unthought-of* situations was used to categorise the 22 recalled incidents, as illustrated by the two episodes summarised below. Table 3.1 shows the distribution of the 22 episodes, according to this classification and according to the physician who related them.

> Potential situation: A patient who never underwent general anaesthesia must be operated on urgently. The intubation is very difficult and the anaesthesiologist cannot put in place the breathing tube. To face, he 'follows the protocol to the letter: several attempts at intubation, chuck, then fast track'. This event is described as 'unexpected' because the anaesthesiologist could not know with certainty in advance that this intubation would be a difficult one: the assessment of risk criteria usually made during the pre-anaesthetic visit, has not been made. But this possibility has been considered: 'it is a situation that every anaesthesiologist feared more than anything in urgency'. 'I had considered the worst'.
>
> Unthought-of situation: At the end of an intervention, when being transferred in the recovery room, the child becomes black, cyanotic and bradycard. The anaesthesiologist begins the resuscitation and calls for help. Two colleagues arrive and take turns to perform cardiac massage. In parallel, the three anaesthesiologists think together in order to understand the event: checking equipment, clinical diagnostics, radiological examination. After 45 min, the diagnosis is made (pneumopericardium). One of the anaesthesiologists performs the technical gesture that will bring the child back to a stable state. The situation is described as 'exceptional'. The physician in charge of the patient did 'not imagine at all that it could happen'. 'I thought that things would proceed as usual'.

Table 3.1 Distribution of the 22 recalled cases of near misses

	Anaesthesiologists (= $A_1, \ldots, _6$)						
	A_1	**A_2**	**A_3**	**A_4**	**A_5**	**A_6**	**Total**
Potential situations	1	3	1	2	2	0	**9**
Unthought-of situations	2	2	2	1	3	3	**13**
Total	**3**	**5**	**3**	**3**	**5**	**3**	**22**

Resilience as the Ability to Define an Envelope of Potential Variability

Nine situations have been categorised as critical situations where the unexpected refers to the time of the occurrence of the event. These *potential situations* are described as relatively frequent ones: the anaesthesiologists '*could watch for it at each* *intervention*'. (In the following, excerpts from the transcript are in *italics*.) In this case, similar stories are told several times by different anaesthesiologists. These situations are always related to the clinical evolutions of patients and do not involve physical or organisational surprise: the events are known events in the field of anaesthesia, learned during training and well described in the literature of the speciality. To handle these situations, care protocols are directly applied. The necessary resources, such as equipment or drugs are available and have been prepared beforehand. In particular, practitioners never called for additional help during the surgical operation: when, in some cases (2/9), the cooperation of two anaesthesiologists was necessary, they worked in pair from the beginning of the intervention. At the extreme, in one case the anticipated risks are such that the two anaesthesiologists decided in agreement with the surgeon, not to perform the operation.

In these nine situations, the adaptability of the system depends on the operators' ability to define an envelope of potential variability before each operation, that is, a set of situations that may occur; in case of which the necessary resources to face are prepared. This envelop is based on the 'objective level of uncertainty' of each event, according to the originality of the event in the field of anaesthesia. For a given operator, the events are more or less uncertain because, according to their degree of novelty, they are more or less known and listed in the field: the relatively common events, such as difficult intubations or spasms, are described in the literature (experts' recommendations, consensus conferences, rules of 'good practices' in services, etc.) making them more predictable, while very rare events, involving, for example, a technical failure and the development of a new clinical pathology are unknown (even unknowable) and very uncertain. For most clinical hazards related, anaesthesiologists mentioned these 'indicators', used to estimate *a priori* the occurrence probability of the event. Some of them are related to diseases: they are codified in the algorithms of care and systematically evaluated before each intervention. But these indexes are not '*always completely reliable*' and their evaluation is closely '*tied to experience*'. Among the scenarios collected, five were *unthought-of situations* when they occurred although these events are identified and known in the field of anaesthesia.

Moreover, other indicators, related to the overall situation of the surgery, are mentioned. For example, some indicators relate to the type of surgery scheduled or to the surgeon who will perform it. Indicators may also include features related to the period of intervention. Finally, in some cases, it is difficult for practitioners to explain the elements that allowed them to anticipate the occurrence of the scenario: they '*felt*' the event or they '*saw it coming*' '*without knowing exactly why*' as in the quote below.

> So I'm going to intubate. And while I have no criteria of difficult intubation and there, it will seem to be paranormal, but 'I have a bad feeling'. I said to myself: 'It stinks. I'll put the child flat to intubation, but I do not know why, it will go wrong'. It's a feeling we have from time to time in medicine. Certainly there are a lot of experiences behind that are probably unconscious.

Therefore, a same event can be a potential or an unthought-of situation according to the operator who is facing it. It depends on the 'capacity [of each one] to project [himself] into future through the current, local, short-term [conditions]' (cf., Chapter 16; Nyssen, 2008). This anticipatory capacity is linked to the experience in the trade, and in particular to the salient situations experienced during practice, as illustrated by this quotation.

> For sure, now, if I have to deal with a lumpectomy I think I will order blood immediately, having been in a situation that was a factor of anxiety, very high anxiety, I will try to avoid generating this anxiety in another situation that seems similar. I will say: 'Last time, if only I had blood'. So I will order blood.

Figure 3.2, adapted from the Leplat's graphic on complexity of tasks in work situations (1988), presents the relation between these two variables: 'uncertainty of the event in the field of anaesthesia' and 'operators' skill to anticipate the future'. It shows two regions corresponding to the two categories of collected episodes.

The region of *potential situations* is where the operator's expectation outweighs the uncertainty of the event. This area corresponds to either well known and well described events for which '*all anaesthesiologists are preparing all the time*', or to less known and more rare events managed by experienced operators whose ability to imagine the future is more efficient.

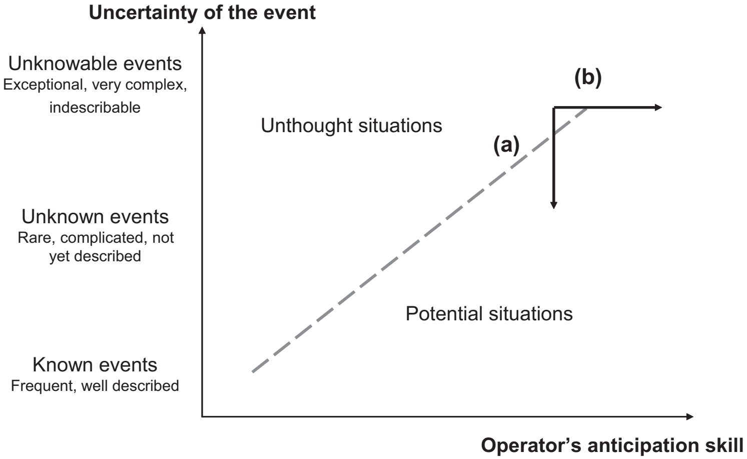

Figure 3.2 Predictability of a given situation for a given individual

The region of *unthought-of situations* where the uncertainty of the event overrides the physician's ability to anticipate the situation. This region corresponds to inexperienced operators or to experienced operators faced with unknown situations.

> The analysis of these potential situations leads to a first approach to resilience: to increase the system's ability to adjust to variability, we must avoid the occurrence of *unthought-of situations*, which means making situations more predictable so that resources can be anticipated beforehand. The graph shows two possible improvement ways.

The first (a) is to increase knowledge about risk situations. To do this, different possibilities are open in the medical field (basic research on diseases including identification of warning signs) and in the field of reliability (modelling types of potential problems and identifying warning signs, through the analysis of past incidents for example (cf., Chapters 7, 15 and 17).

The second (b) is to develop the operator's ability to project themselves into future in real conditions (cf., Chapter 16). This is the path of vocational training and especially of training based on the analysis of practices and on reflective activities (Falzon, 2005; Mollo and Falzon, 2004) that may aim for example, at 'clarifying the foreboding' or highlighting 'embedded' expert's knowledge.

Resilience as the Ability to Diagnose that the System Leaves the Envelope of Potential Variability

Thirteen of the 22 cases collected are situations that were not envisaged by the anaesthesia team in charge of the patient before their occurrence. In the recalled stories, these situations are described as exceptional. Indeed, the 13 situations related are unique: there is no similarity between the different cases scenarios. But, as mentioned previously, in five cases, the *unthought-of situation*, which surprised the practitioner during the surgical intervention, concerned solely the patient's clinical course. And, ultimately, it proved to be a known pathology, described in the literature and for which procedures are defined (e.g., pneumopericardium). The other cases (eight) are situations that involve an organisational or technical unexpected event, like breakdown of equipment or problems of cooperation with the surgical team, as illustrated by the case summarised below.

> The child who must undergo surgery presents a latex allergy. The anaesthesiologist informed of this fact, put all the measures in place to prevent the occurrence of this allergy: the operation is scheduled early in the planning, all materials containing latex are removed from the operating room, the whole team is informed. After the patient asleep, the anaesthesiologist leaves the room to care for another surgical operation and lets the anaesthesiologist resident monitoring and controlling the surgery of the allergic child. He is recalled after a few minutes because the child has an anaphylactic shock (a severe allergic reaction). The possibility of an allergic reaction to latex has been anticipated by the anaesthesiologist but it is unthinkable for him that there is latex in contact with the child, since he has made every effort to remove all the latex from the operating room. The fact that the resident has picked up a protection sensor in latex outside the room is, at this moment, an unimaginable event. To cope with the situation, the anaesthesiologist calls a colleague for help, who quickly detects the latex in the child's nose.

Although the situation has not been considered before its occurrence, it was sometime (2 cases out of 13) immediately understood and managed through the strict application of an existing protocol. But in the most cases of *unthought-of situations*, the strict application of existing protocols cannot directly deal with the event (11 stories out of 13). In these cases, anaesthesiologists had not envisaged such a situation and were in difficulty.

Identifying and Understanding the Situation

The diagnosis is not immediate, the sense of what is happening is lost (Weick, 1993), and since the problem is not identified, it is impossible to implement a protocol and to bring the situation with certainty in stable condition. The anaesthesiologist is here confronted with the cognitive trade-off for managing dynamic situations (Amalberti, 1996): he must choose between understanding – that is, maintaining an unstable survival state of the patient to pursue his reasoning and establish a correct diagnosis before acting – and agreeing not to understand that is, choosing one of the possible protocols, according to the hypotheses of diagnostic.

Implement Ways to Cope

The situation is understood but protocols cannot be applied because provided technical gestures do not work or surgical teams do not meet the demands of the anaesthesiologist. The anaesthesiologists may then either continue to apply an inefficient protocol provided by the organisation, or he may conduct actions 'beyond protocols'.

In both cases, physicians must take decision(s) based on the assessment of the risk/benefits ratio of the different options, generally under very high time constraints. Based on these decisions, additional resources are mobilised according to the specificity of each situation: medical examinations (such as radiology), blood order for transfusion, use of emergency drugs etc. In particular, we note that in nine of these 13 *unthought-of situations*, the anaesthesiologists called on colleagues for help. Thus, calling for help appears as a transversal strategy for the different *unthought-of* scenarios collected. Earlier results (Cuvelier and Falzon, 2008) have shown that this call for help is not simply related to the identification of a need (estimation by the practitioners that the resources provided inside the envelope of potential variability are insufficient) but is rather a trade-off decision between multiple criteria: the availability of colleagues, the time of the day, the function of the call etc. These decisions to mobilise additional resources (and in particular to call a colleague) are the observable signs of the shift of the system from a 'thought-of' state considered *a priori* as potential, to an 'unthought-of' state. In this sense, they have been described as 'pivotal decisions.' Resilience in these *unthought-of situations* lies thus in the operator's ability, not only to detect, but also to decide that the system leaves the envelope of potential variability.

Enhancing Resilience: Paths for Progress

This study highlights the role of two crucial decisions in the management of uncertainty by anaesthesiologists. The identification of these two decisions provides opportunities for research and actions to improve resilience in the practice of anaesthesia.

The first decision is the preoperative definition of an envelope of potential variability of the surgical intervention. Results show that the construction of this envelope does not only depend on the situation's level of uncertainty (frequency of occurrence, objective prediction criteria, etc.), but also on the ability to project oneself into the future, an ability which is closely tied to experience. In this perspective, an investigation of the way in which anaesthesiologists construct the envelope of potential variability, taking into account the role the experience in this construction, would be helpful. One possible avenue for improving the resilience could be that of training through reflective practice. Moreover, the definition of the 'envelope of potential variability' may also be questioned at various organisational levels in the system and between different trades: is, for example, the envelope of potential variability defined by the anaesthesiologist, consistent with the one defined by the surgeon or with the one defined by the management staff?

The second decision concerns the occurrence, during the operation, of an *unthought-of situation* which trespasses the envelope initially defined. In such a situation, anaesthesiologists may or may not understand the problem. But in all these cases, they must decide if the problem can be managed with the resources provided in the envelope, or if the mobilisation of additional resources, including the call of colleagues, is necessary.

Another area for further study is to investigate the mechanisms of this decision on the spot, using simulators and in depth recall interviews. It consists not only in identifying the factors involved in this decision (severity of the situation, changing dynamics of the process, availability of resources etc.) but also in understanding how the practitioner evaluates their own capacity to manage the event. Once again, the role of experience can be questioned, wondering, for example, if the decision taken by a novice to call a colleague in an *unthought-of situation* (that is unpredictable for this person in this context, so not necessarily very complicated) is based on the same mechanisms as the decision to call for help made by an expert in an unknown situation.

Acknowledgements

This study is supported by a grant from the 'Haute Autorité de Santé' (France). We thank the anaesthesiologists who agreed to participate in this study. We also wish to thank Professor Leplat for his availability and his interest in our work.

This page has been left blank intentionally

Chapter 4
Training Organisational Resilience in Escalating Situations

Johan Bergström, Nicklas Dahlström, Sidney Dekker and Kurt Petersen

Over-reliance in high-risk industries on prescriptive emergency procedures and the capacity of high-fidelity simulation has initiated the search for new steps for training for unexpected and escalating situations. This chapter outlines a theoretical framework describing the adaptive and flexible competencies that add up to an organisation's resilience in escalating situations. To make the framework practically useful the chapter also presents guidelines for scenario design aimed at the training of generic competencies in unexpected and escalating situations. We finally show the potential of using training programmes, adhering to the given scenario guidelines, for securing front-end resilience in unexpected and escalating situations by presenting an experiment that was performed with Swedish Fire Safety Engineers.

Introduction

This chapter focuses on the ability to maintain organisational resilience in unexpected and escalating situations. The concept of escalation is described by Woods and Patterson (2000) as a dynamic process in which an initial irregularity develops into a continually deteriorating situation and starts affecting other areas in an accelerating tempo, with consequences that are difficult to overview and impossible to predict (for a discussion about unexpectedness, see Chapter 3 in this volume).

In a previous book on Resilience Engineering, Dekker et al. (2008) proposed that training aimed at the development of generic and non-domain specific competencies can increase organisational resilience in unexpected and escalating situations. One example of this being that the enhancement of adaptive and flexible competencies can support problem solving in a group by enabling group members to disconnect from prescribed role behaviours and routines.

Building further on their argument this chapter outlines a theoretical framework describing the adaptive and flexible competencies that add up to an organisation's resilience in escalating situations. The framework is called *Generic Competencies in Management of Escalating Situations*. The term generic is chosen to distinguish the competencies from the domain-specific, and more technical, competencies that are usually practised in domain-specific training, as in use of high-fidelity simulations. This chapter will also present guidelines for scenario design aimed at training of generic competencies in unexpected and escalating situations. We will also show the potential of using training programs, adhering to the given scenario guidelines, for securing front end resilience in unexpected and escalating situations by presenting an experiment that was performed with Swedish fire safety engineers.

Generic Competencies in Management of Unexpected and Escalating Situations

In order to explore and explain the competencies which can support organisational resilience in unexpected and escalating situations, a theoretical framework has been developed. The framework can be used as an explanatory tool when discussing the concept of organisational resilience, as well as a tool for qualitative evaluation of decisions and actions in unexpected and escalating situations. The framework has been developed through case studies of unexpected and escalating situations in various industries as well as studies of teams managing escalating situations in simulated environments. The competencies should be regarded as team competencies and not the competencies of an individual team member. Figure 4.1 shows a simple illustration of the framework.

The first competence category is *Information Management*. In terms of data processing the effects of the escalation is dual. With the escalation follows an increased amount of data to process (Woods and Patterson, 2000). On the one hand having greater access to data may be beneficial in principle, but on the other the flood of data challenges the ability to find what is informative and meaningful in the data flow. Woods et al. (2002) refer to this as the data availability paradox. One way to make sense of the high incoming data load is to use shared and explicit goals (Dörner, 1996), based on which the incoming data can be sorted, distributed, and shared.

In the *Communication and Coordination* processes the importance of knowing each other's roles and tasks in the team is emphasised (Klein et al., 2005). There are, however, two main challenges facing the team structure during an escalation. The first challenge is to recognise how the demands on each participant develops during the escalation and to respond to possibly overloaded team members by reformulating certain tasks or bringing in additional resources. The second challenge is to be aware that escalating situations do not necessarily adapt to our (often rather bureaucratic) idea of how the respond system should be set up

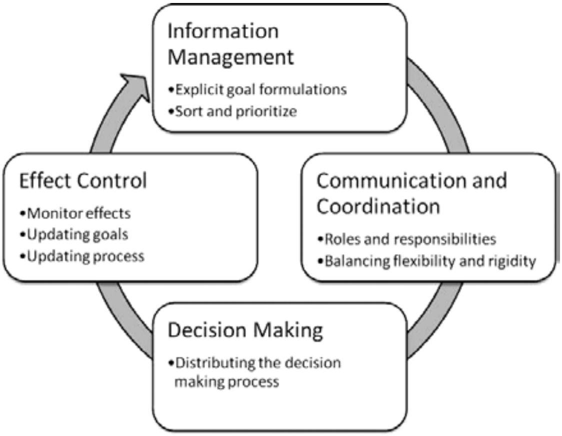
Information Management
•Explicit goal formulations
•Sort and prioritize
Communication and Coordination
•Roles and responsibilities
•Balancing flexibility and rigidity
Decision Making
•Distributing the decision making process
Effect Control
•Monitor effects
•Updating goals
•Updating process

Figure 4.1 A framework of generic competencies in managing escalating situations

(Uhr, 2009). There is therefore a need for constant monitoring of whether the response organisation is suited to manage the situation at hand and adapting the organisation's role- and task structures as the situation develops (Brehmer, 2008). The organisation has to be able to balance between a predefined role structure and a flexible structure to respond to the dynamics of the escalation (Heath, 1998).

The third competence category is named *Decision Making* and focuses on the strategies used to make decisions in escalating situations. Decisions in escalating situations cannot be based on consensus in the managing team. Such a process would simply 'drown' the team members in too much data to process, and result in reactive decision-making behaviour. Nor is a hierarchical structure, with a team leader making all decisions, likely to be successful, because of the workload that would face such a team leader responsible for making all decisions (Bergström et al., 2009). Instead the shared and explicit goals need to be used for decision making in a distributed decision-making environment. In such a process decisions are made by all participants. Such a strategy poses high demands on the information sharing strategies in the team. Information about decisions made, and updates of goals, needs to be shared in order to keep such a distributed decision making environment functioning.

In the fourth category; *Effect Control*, the importance of constantly monitoring and updating the process, by which the escalating situation is managed, is stressed. Goals have to be questioned and updated, task descriptions and areas of responsibility have to be negotiated and adapted to the dynamics of the situation and questions like 'what could be wrong in our understanding of this situation?' needs to be raised.

Scenario Design

Scenarios developed to practice the generic competencies in escalating situations need to comprise some basic principles of escalating situations, as described by Woods and Patterson (2000).

- There is a cascade of effects in the monitored process.
- The cascade of effects should demand an increase of cognitive activities among the participants.
- The nature of the cascade of effects should demand an increase in coordination among the participants. The process cascading should therefore not be isolated to one particular participant's area of responsibility, but instead demand different reactions by all participants.
- The cascade and escalation should be a dynamic process.

In addition, in order to maximally utilise the escalation scenario and enhance people's generic abilities to manage it, the scenario should do the following.

- Try to force people beyond their learned roles and routines. The scenario can contain problems that are not solvable within those roles or routines, and forces people to step out of those roles and routines.
- Contain a number of hidden goals, at various times during the scenario, that people could pursue (e.g., different ways of escaping the situation or de-escalating it), but that they have to vocalise and articulate in order to begin to achieve them (as they cannot do so by themselves).
- Include potential actions of which the consequences are both important and difficult to foresee (and that might significantly influence people's ability to control the problem in the near future). This can force people into pro-active thinking and articulation of their expectations of what might happen.
- Be able to trap people in locking onto one solution that everybody is fixedly working towards. This can be done by garden-pathing; making the escalating problem look initially (with strong cues) like something the crew is already familiar with, but then letting it depart (with much weaker cues) to see whether the crew is caught on the garden path and lets the situation escalate.

- Or the scenario, by creating so much cognitive noise in terms of new warnings and events, should be able to trip people into thematic vagabonding – the tendency to redirect attention and change diagnosis with each incoming data piece, which results in a fragmentation of problem-solving.

Training Generic Competencies

To evaluate the possibilities to build front-end organisational resilience in unexpected and escalating situations, by using training programmes adhering to the scenario design criteria above, an experiment was made at the Swedish Civil Contingencies Agency's (at the time of the experiment known as the Swedish Rescue Service Agency) school in Revinge, Sweden. A two-day crisis simulation exercise was used to practice the generic competencies of two experimental groups.

The simulation allowed five to seven participants to assume different roles on the bridge of a poorly maintained passenger vessel caught in a stormy night on the Atlantic Ocean (Strohschneider and Gerdes, 2004). The simulation programme adhered closely to the principles for scenario design, listed above. During the simulation different events occurred (see Figure 4.2 for a simple sketch of the first day scenario) that increasingly demanded that the participants established strategies to apply generic competencies to prevent the situation from escalating beyond their control.

The simulation was complex and was run by a computer. The participants were able to effect the development of the situation by the actions they took. The simulation provided information to the participants in the form of computer printouts. Beyond blueprints and maps there was no visualisation of the simulation and the participants were not equipped with any predefined strategies for managing upcoming situations. The simulation was part of a two-day training programme which, apart from one simulation session (taking about three hours) also included lectures, discussions and debriefing sessions.

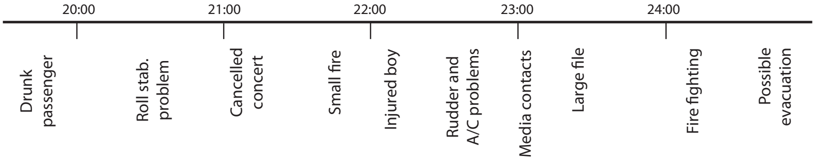

Figure 4.2 Simple outline of the first day scenario. Other things might occur depending on the actions taken by the participants along the session

From a course with 23 Fire Safety Engineers, on a year-long training programme to become incident commanders in rescue services, half received the two-day programme before scheduled emergency management staff training on their course. At this stage of their programme the Fire Safety Engineers had already completed six months of rescue operation training in a rescue service-environment. They had practised standard emergency response in mostly fires and road accidents, but also in more complex scenarios like chemical accidents.

During the simulations data collections were made with regard to the students' abilities to use generic competencies. Data collections were later made during the Fire Safety Engineers' emergency management staff training, that is, where they were performing within the boundaries of their own domain (playing the role of the emergency staff responsible for controlling and coordinating the various actors in a particular emergency operation). That training was given two weeks after the first, and one week after the second, experimental group went through the non-domain specific training outlined above. Each of the four groups (the experimental groups being intact from the prior non-domain specific training) was individually observed when respectively managing two different emergency scenarios taking approximately one hour each. In these scenarios the teams were situated in a room from which they communicated with other actors (e.g., rescue services on the scene, the alarm central, and anyone else that they wanted to get in contact with) over phones and radio devices. The facilitators running the scenario, and playing all the roles that the participants would like to get in contact with, were situated in another room. Prior to going through the emergency management team training all groups participated in a brief introduction lecture in emergency staff structure (role of the staff and roles in the staff). Differences in the use of generic competencies between those engineers who had received the two-day program and those who had not, were observed and analysed.

In analysing the performance of the teams, a distinction was made between the two perspectives of *process* and *outcome*. Outcome relates to the quantitative results of the simulation; primarily numbers of injuries and casualties, and damage to the ship (in the case of running the vessel bridge simulation). As a performance measure, however, the outcome of the simulation is dependent on the interaction between the participants and the facilitators of the simulation and therefore renders it a less reliable measure of performance. *Process* refers to qualitative aspects of how the group managed situations encountered in the simulator sessions in relation to accepted and recommended practices for emergency management. When analysing team performance based on the generic competencies in managing escalating situations, we were interested in the qualitative aspects of the *process*.

The two experimental groups were initially unsuccessful in handling the simulated vessel during the non-domain specific simulation program. None of the groups established strategies to handle the high data load, they did not state explicit goals and they held rigidly to their assigned roles and tasks without updating the process to the dynamic nature of the situation. Shortly into the simulation the teams' performance could be described as normal operations behaviour, focusing on what to do to solve current problems based on their urgency, rather than trying to refocus on how to create structures and strategies to solve problems based on an assessment of their importance. However both teams did improve their performances at the second day's exercise and expressed that the exercise had been useful for their training.

During the following emergency management team training an increased ability to apply generic competencies was demonstrated by the experimental groups, compared to the control groups that had only received the ordinary emergency operation training.

Both experimental groups chose a goal-driven *information management* process in which participants with the assigned roles of moderators controlled the information flow into and out from the team as well as within the team. One of the teams used a strategy in which the moderator wrote important information on notes that were then distributed to the participant responsible for the particular task of interest. In the high load of incoming information, both experimental groups used explicit goal formulations to guide the process of determining what information that was relevant. The control groups did not use such a strategy, but sorted incoming information based on perceived urgency rather than on importance. In none of the control groups was a particular role assigned to manage the information-sharing process. Rather, the one closest to a calling phone would take the call and then update everyone in the team of the content of the incoming information. In one of the control groups the assigned team leader was caught up in talking on the phone during an entire exercise session.

The *communication and coordination* process was characterised in both experimental groups by the participants being assigned specific roles. This was the same with the control groups. However, when managing the dynamic scenarios, the role structures of the control groups immediately broke down and the process of assigning specific tasks to specific participants seemed unorganised and *ad hoc*. In both experimental groups the roles of the assigned team leader was to overview the process and suggest updates of goals, role formulations and tasks. The role of the team leader in the control groups was not as clear and several times resulted in the team leader being bogged down in telephone and radio communication. In the experimental groups the different tasks to manage were distributed among the different participants and they regularly gathered for briefing each other on the latest decisions made and other information that they thought was needed for other team members, based on the shared goals. The control groups also gathered for briefings, but without a clear role structure these briefings were held as soon as any new information was received so that everyone could be updated on all information. With the high load of incoming information this meant the control groups were spending more time updating each other with all available information than actually making judgements and decisions.

The lack of explicit goals and the broken role structure made the control groups establish a *decision making* process in which decisions were made in consensus. The experimental groups instead used the explicit goals and role formulations to establish a distributed decision-making process in which decisions were constantly made by all team members, based on their specific tasks to manage. They then used the briefing sessions to brief each other of decisions made and agree upon new goals, rather than using the briefing sessions for actually making decisions. The briefing sessions in the experimental groups were in that way proactive, setting the goals for the decisions to come, while the briefing sessions in the control groups were reactive with a focus on making decisions based on what had already happened. Finding this proactive switch seemed to be of decisive importance for the organisational resilience in escalating situations.

In the dynamic environments that were simulated and the rather short exercises that were conducted, none of the teams felt that they had the time to establish an *effect control* process in the sense of following up decisions made. However the experimental groups used the briefing sessions to update the agreed-upon goals and revise role descriptions based on the development of the situation and the workload on different members of the team.

The ability to establish strategies based on their generic competencies made the experimental groups able to establish more proactive processes, focusing on expectations rather than history, than the control groups, which rather were stuck in the inability to sort, prioritise and distribute information and tasks. A summary of the observations made during the staff exercises is shown in Table 4.1.

The most significant difference between the groups was however observed during the debriefing sessions. During these the control groups, often taking a defensive position, expressed confidence in the roles and procedures that they would use in the 'real world' in case of a 'real world crisis'. They questioned the usability of the staff exercises because of this lack of correspondence with the 'real world' situation. After receiving the non-domain specific training the experimental groups instead expressed the belief that there is no need for exact coherence with the 'reality' to gain learning from an experience. In this they showed a deepened knowledge about the unpredictable nature of the escalating crisis, and an understanding of the need to adapt a crisis management system to the dynamics of the escalating situation at hand.

Table 4.1 A comparison between the groups that had received the simulation training program and the control group's performances at the staff exercises

The experimental groups	The control groups
Indistinct roles at high data load	Hardly any roles
Using goals to establish a proactive process	Reactive process
Briefing sessions to update each other of decisions made and revise the process	Briefing sessions to establish a consensus on what decisions to make
Team-leader over-viewing the process	Team-leader often stuck up in radio-communication and operational decision-making
Assigned participants the task to sort and distribute incoming information	Who answers the phone was selected by chance
Tasks were performed	Thematic vagabonding
Some explicit goals	No explicit goals

The experimental groups also made far more qualified analyses of their performance and their shortcomings. Their statements showed an understanding of the need for generic competencies and the difficulties in establishing strategies for applying them. A summary of the observations made during the debriefing sessions is shown in Table 4.2.

Discussion

The experiment outlined here showed that designing a training environment in which people really faced the uncertainty an unpredictability of escalating situations (whatever the domain) generated 'resilient' competencies that were not generated by current training strategies aiming at drilling correct behaviour in known scenarios.

Table 4.2 Differences in reasoning at the debriefing sessions after the staff exercises.

The experimental groups	The control groups
Identifies the problems in doing other peoples' work	No understanding for the importance of roles
Discusses the difficulties in formulating explicit goals and the benefits from doing so	Believes implicit goals are capable of guiding the management
Discusses the difficulties in being proactive	Wrongly believes that some actions were proactive
Generally good in evaluating their own actions	Express believes that in real life there are predetermined roles and procedures for all situations

It also showed that the non-domain specific training deepened the understanding of the nature of escalating situations and the difficulties in managing them. The potential is great to apply this sort of training in various industries that demand rapid and well structured response to escalating situations, although more research and further testing is needed. However, as a part of the overall aim to increase the understanding of resilient organisations and their characteristics, non-domain specific simulation have already proved to be an effective tool.

It also seems as if the training of generic competencies benefited from taking place outside of the participants known domain. This was shown by the control groups' defensive attitude in the debriefing sessions in which they criticised the training for not being an exact replica of 'the real world' (thereby suggesting that the entire training was useless). Having the training taking place in a, for the participants, completely unknown domain seemed to give an understanding of the unexpected nature even of ‘the real world,’ and it also seemed to remove the prestige-loss that follows from failing in the participants everyday domain. That can be seen by participants often discussing failure in an unknown domain, not by taking a defensive position but by more complex discussion about the difficulties of managing any unexpected and escalating situation.

The theoretical framework explaining generic competencies for proactive crisis management proved to be a useful tool for contextualising statements made, and strategies chosen, by the observed teams in their managing of the unexpected and escalating situations. The framework should, together with the scenario guidelines, also be a useful tool for development of new programs and methods for managing and training of escalating situations in numerous industries like aviation, ship management, health care and the nuclear industry.

The generic competencies in management of unexpected and escalating situations are of decisive importance for any organisation to secure resilience in unexpected and escalating situations. These competencies must be practised, not by drilling prescriptive plans and procedures, but by adhering to the principles of the very nature of unexpected and escalating situations. Developing scenarios and training programmes based on the guidelines presented here could be one way of starting such a process.

This page has been left blank intentionally

PART II
Dealing with the Critical

This page has been left blank intentionally

Chapter 5
Monitoring – A Critical Ability in Resilience Engineering

John Wreathall

The prologue to this book outlines the broad defining characteristics – the four cornerstones – of Resilience Engineering. One of these is monitoring. Every organisation concerned with safety has one or more metrics that are used to judge whether the levels of safety in the organisation are acceptable. Using a common definition of safety, the organisation needs to know if it is 'free from unacceptable risk'. This metric is often the number or rate of accidents or injuries (or deaths) over some period of time or the time between events as shown by Figure 5.1.

While such measures may provide some assurance that safety has not been entirely out of control in the past, they are of little use when considering how to manage safety in the future and could even be harmful. First, top-level outcomes such as accidents or serious injuries have a large component of chance in their occurrence. Second, the events do not, themselves, provide information about causes and fixes. Finally, when the measures show a period of good performance, complacency starts to set in. Indeed, when a period of successful performance approaches a record, there can be tremendous pressure for people not to report possibly marginal occurrences and somehow 'lose the game'.

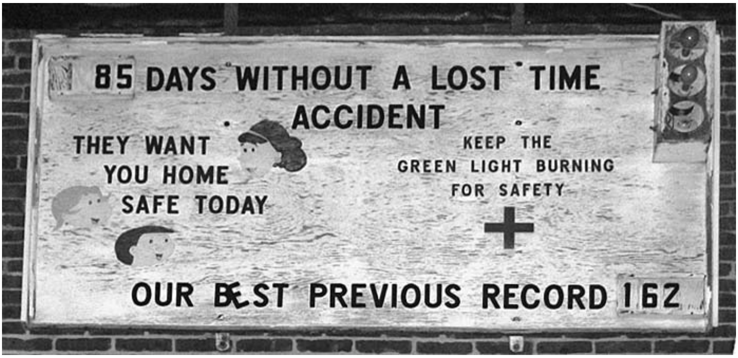
85 DAYS WITHOUT A LOST TIME ACCIDENT
THEY WANT YOU HOME SAFE TODAY
KEEP THE GREEN LIGHT BURNING FOR SAFETY
OUR BEST PREVIOUS RECORD 162

Figure 5.1 Typical industrial safety sign (Photo © Sheena Chi, 2009)

But measures like that shown in Figure 5.1 are of little or no use in preparing for foreseen and unforeseen adversity or in managing the proactive processes by which safe and efficient performance is achieved. Since the environment of the organisation and its own internal processes are both dynamic, last year's safety performance (or last month's or yesterday's) are at best weak indications of how today's and tomorrow's will be.

The Role of Indicators in Measurement

Measurement of the processes is an essential part of any organisation. The phrase, 'You can't manage what you don't measure' is an old management adage and appears in more websites that this author could count. However, in developing Resilience Engineering, the role of indicators is on the one hand crucial and on the other also an area that is underdeveloped relative to other aspects of measurement.

While there are many metaphors for how resilience applies to the management of organisations, one that is useful in the context of measures and indicators is the classical control theory model shown in Figure 5.2. The activities in the organisation are represented in the box labelled 'processes'. These processes comprise the main activities in the organisation, such as production (in its various forms), coordination, finance, etc. They accomplish a variety of outputs that include physical production, safety, economic performance, etc. In the classical control model, these outputs are evaluated by means of a 'Model of the Processes' that via a controller adjusts the inputs to the processes guided by the requirements. This set of activities takes place within an environment (the public community, regulatory bodies, the financial and business environment, etc.). The requirements may be set within the organisation itself or by the environment (such as public safety goals, financial or work demands). Some of the outputs likewise go to the environment, like work products, financial matters (including taxes and interest on loans), etc., while some are internal (record keeping, retained profits, etc.).

The management of safety has traditionally followed the classic control theory model. Changes in safety practices and policies usually do not occur unless there is a change in the 'safety output' of the process – a loss of life, a serious injury, an explosion or crash or some other serious event. Then the management reviews the event through using its 'model' of the processes, makes some changes (putting in place a new safety policy or a new barrier perhaps) and let the business continue. The sign shown in Figure 5.1, for instance, reflects the thinking that simply making the output more visible to the workforce will change the processes by reminding them that worker safety is an outcome.

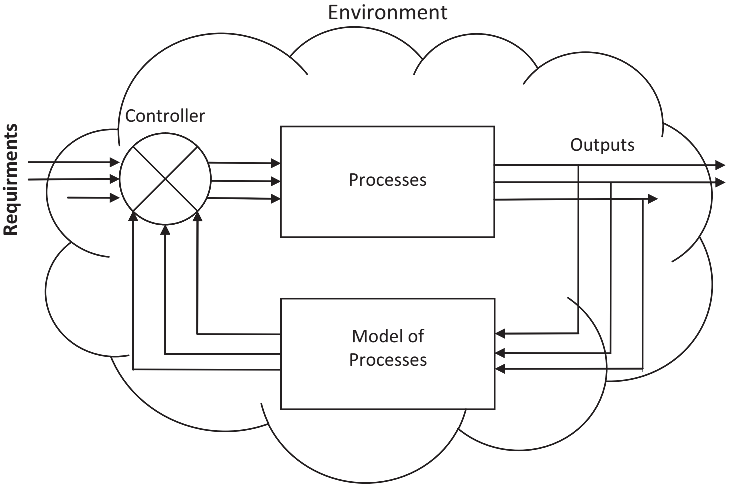

Figure 5.2 Classic control theory model

Resilience Engineering emphasises the need to be proactive in the management of core processes, including but not limited to safety, and to anticipate (and hopefully forestall) major changes in safety and other critical performance domains. Simply relying on outcomes as implied in Figure 5.2 will not suffice. In order to be proactive, more information is required, as shown in Figure 5.3. In this version, data are gathered not just from the outputs of the processes but from intermediate activities along the way, as shown by some of the dashed lines from the processes to the model. These data we will refer to as indicators. Their purpose is to provide information about what is happening in intermediate stages of the processes before outcomes change significantly, so that management can take actions to forestall the adverse outcomes.

Indicators can also be developed for the changes in the environment that may impact the system. Obvious examples would be to sense upcoming changes in the financial environment that may affect the processes (e.g., changes in the cost of raw materials, changes in customer requirements, a tightening of credit). Anticipating changes in the environment and making anticipatory changes in the processes to accommodate them are an important part of the long-term stability of any organisation.

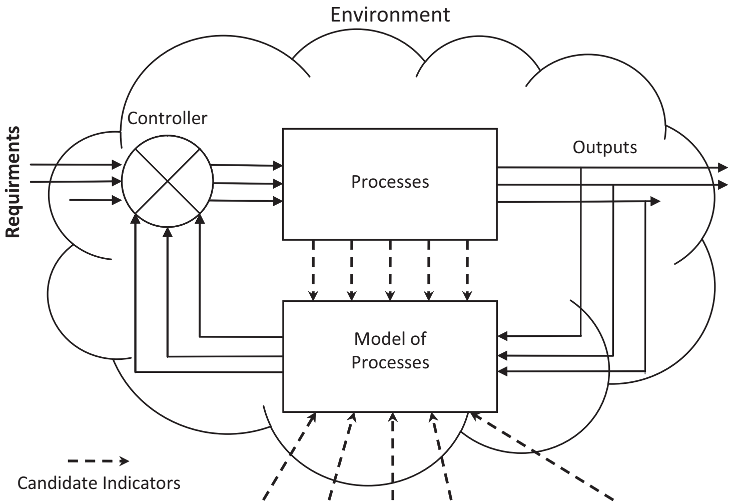

Figure 5.3 Control theory model with indicators

Westrum (1999) has pointed out that using 'faint signals' often is a critical feature of resilient organisations. Faint signals are early indications of problems that start to occur in a project or hints of coming trouble in a process. One example could be an increased number of questioning-type calls from a customer about how things are going. Nothing explicit, but often after a problem has been uncovered, they would be recognised as early warnings.

Selection and Basis for Indicators

Wreathall (2009) proposed that 'indicators are proxy measures for items identified as important in the underlying model(s) of safety. As such they are uncertain and often only distantly connected to idealised measures that rarely are available in practice'.

In practice, it is rare to find explicit models that provide a formal basis for identifying measures. However, knowledge often exists about relationships that are important to safety and that can be used to create working models. In the study of fatigue (Chapter 6), a set of relationships between safety, fatigue and its underlying mechanisms was articulated (if not formally presented as a model of safety) and those relationships were used to create metrics for monitoring the risks from fatigue. In their paper, the authors acknowledge that it is not yet possible to create a fully developed model but that sufficient information exists to create a useful basis for management of the processes of concern.

In the study of cognitive strategies in air traffic control (Chapter 8), the T^{2}EAM performance model was developed based on existing research related to cognitive strategies in similar kinds of tasks, though often in other industries. This performance model provided the basis for developing ways to measure.

As a practical matter, several studies have used *ad hoc* methods for selecting indicators. As described in Chapter 7, the critical factors associated with an outage indicating a potential to be problematic were identified through brainstorming workshops with experienced project managers and others often involved in supporting projects when they get into trouble. In other words, even though these were based on learning from events in the past, they were considered to be sufficiently robust that they would remain effective for the next cycles of outages (1–2 years), at least from the perspective of testing the concept. Similarly in work described by Wreathall (2006), the selection of indicators was based on factors most frequently associated with problems in human performance issues in the nuclear industry's collective experience.

One point worthy of note in the experiences referred to above is that while the initial search for factors related to performance began by looking for factors that contributed to problems in performance, the development of these measures uncovered factors that were critical to the success of the system. For example, in the system reported by Wreathall (2006), the data gathered for particular parameters indicated a consistent positive pattern. When discussions were held with the reporting staff, they indicated that these were important factors to ensure successful performance.

Nature of Indicators

Wreathall and Jones (2000) proposed a set of desired characteristics to guide the selection of indicators. Preferred indicators are as follows (in decreasing order of priority):

- Objective: They are based on observable and non-manipulatable sources.
- Quantitative: They are measurable and can identify when changes in performance occur.
- Available: They can be obtained from existing data.
- Simple to understand/represent worthy goals/possess face validity: It has been found in the author's experience that once indicators have been established, they become pursued as goals in themselves. If the indicators in themselves represent something worthy, this pursuit helps the organisation's performance.
- Related to/compatible with other programs. In any modern enterprise, data abound in excess. It is generally undesirable to add an additional data generation program to existing activities.

The experience to date shows that there are a variety of different kinds of measures. The examples of Chapter 6, related to the Fatigue Risk Index model principally involve measures directly associated with objective items such as duties starting early in the morning (before 06:00), working periods of greater than 5 days and the numbers of flights. In the T^{2}EAM model (Chapter 8), the measures are scores associated with the styles of behaviour used by the controllers as observed by the various subject matter experts. In Chapter 7, a variety of different measures are used; some are objective scores like the amount of emergent (unplanned) work, prolonged duration of work, personnel turnover and the existence of special conditions (extreme weather, occurrence of major holidays, etc.). Others include more subjective judgements like a drop in moral, and a reduction in the quality (as well as quantity) of site-to-HQ communications such as perfunctory status reports. Thus, within these three example studies, almost the entire spectrum of possible measures is seen as providing indicators.

Leading and Lagging Indicators

In discussions about using indicators, much is made of whether particular indicators are leading or lagging, almost as if this is a theological issue. The basic issue of which indicators are leading or lagging is a pragmatic one. Consider the arrangement shown in Figure 5.3. The indicators are called leading if they provide information such that control actions can be taken in time to forestall an unacceptable change in one or more of the core outputs (e.g., safety-, financial- or quality-related), or at least the management can anticipate and mitigate the adverse changes – core behaviours of resilient organisations. Indicators are called lagging if they reflect changes in core outputs that have occurred. The parameter shown in the sign in Figure 5.1 is an example of a lagging indicator. It reflects the past performance in managing safety.

However, the issue is a little more complex when there are multiple time-scales and levels for control actions. Sometimes actions may be taken quickly at a local level. For example, increasing the staffing temporarily in a production department to cope safely and reliably with a sudden upsurge in demand for a product may be decided within hours or days of an indication of the upcoming demand. However, there may be a need to change the total employment in the company or redesign the processes to accommodate extended periods of new demands. A lagging indicator related to the short-term staffing change (e.g., numbers of workers suffering from fatigue-related injuries) could act as a leading indicator of the need for systemic changes. Thus the terms leading and lagging apply both to the type and level (local or system-wide for example) of control action and the timing of the actions.

This classification is truly pragmatic and many happy hours can be spent debating with colleagues much deeper interpretations. Recently a special issue of Safety Science was dedicated to a discussion of indicators and part of which was preoccupied with just this question (e.g., Hopkins, 2009a, 2009b; Hale, 2009; Wreathall, 2009 among others). The purpose in each case is, of course, to allow the management of the system or organisation to take actions that prevent adverse outcomes in the event of challenges.

Chapter 6

From Flight Time Limitations to Fatigue Risk Management Systems – A Way Toward Resilience

P. Cabon, S. Deharvengt, I. Berechet, J.Y. Grau, N. Maille and R. Mollard

Because of the development of the 24/7 operations in various industries, human fatigue is today considered as one of the major risks for safety. To date, the prescriptive approach through the regulation of duty hours is the traditional way to prevent fatigue. However, besides the inherent rigidity of regulations from an operational point of view, this often fails to take into account all of the complex dimensions of fatigue. In order to cope with this complexity, Fatigue Risk Management Systems (FRMS) are progressively emerging. Rather than setting absolute duty time limitations, a FRMS approach evaluates each operation in terms of fatigue risk. FRMS can be seen as a concrete way to engineer resilience because it requires the organisation to adjust its functioning by re-introducing safety managed by humans in addition to safety by regulations. This chapter presents a concrete application of FRMS in civil aviation. The whole process of the FRMS is described, from the use of predictive models of fatigue to minimize the risk at the aircrew scheduling stage to the development of fatigue related indicators. The principles of the FRMS are discussed from the Resilience Engineering perspective.

Introduction

The regulatory developments for flight and duty time limitations (FTL) in aviation are the results of different processes. The technological and operational developments of aviation (e.g., aircraft range capacities, opening of new routes), the changes in economic and social context (e.g., shortage of pilots, salary negotiations), and implementation of safety management strategies (e.g., training, oversight) are factors that contribute to establish different forms of consensus among states but also between various airlines. In short, when the airline industry struggles in a fiercely competitive environment, it needs operational flexibility which can be gained through flexible usage of the pilot workforce. This is realised in exchange for various compensations in a win–win situation. However this complicates the management of safety since prescriptive regulations are not adapted to achieve such flexibility.

Amalberti (2006) suggested two reasons, external and internal, for a system to change its level of resilience, sacrificing its competitiveness for a more acceptable safety level. In our case this was the Member States – a European Parliament decision in 2006 to implement in 2008 a Europe-wide prescriptive regulation on FTL (i.e., EU-OPS, subpart Q). Based on the ability offered by the legal text to maintain member states national schemes for specific boundary conditions (reduced rests and split duty), it was proposed to engineer a new resilience level based on FRMS. This construction was developed with lessons learnt from past experience in engineering human factor solutions into aviation (Deharvengt, 2007): scientific expertise and a long term process of implementation from the regulator perspective. A large scale effort was put into place by Direction générale de l'Aviation civile (DGAC) to support a scientific team to develop guidelines for airlines and regulator for the new regulation.

Several industries and airlines have already evolved towards a non-prescriptive approach focusing on fatigue risk management rather than solely on the compliance to a FTL scheme. In aviation, New Zealand has the longest experience in the development of FRMS. In 1995, the regulations were altered so that air operators could either comply with a standard prescriptive scheme or apply an alternative, approved company-specific scheme (Civil Aviation Authority of New Zealand, 2007). In this last case the operator has to take into account additional factors that may result in fatigue (Signal et al., 2008; e.g., rest prior duty, effects of time zone change). The introduction of Ultra Long Range flight by Singapore Airlines in 2003 is another example of a FRMS application. The Civil Aviation Authority (CAA) of Singapore has allowed the airline to operate those flights using the results and recommendations from a scientific study based on biomathematical modelling (Spencer et al., 2002; Spencer and Robertson, 2007). In Europe, easyJet became the first major airline to be granted alleviation from the current FTL in 2005 (Stewart 2007). The UK CAA agreed the alleviation based on the results of a safety case report of a six month roster trial.

As fatigue differs from many other workplace hazards because it is impacted by all waking activities, not only those that are work related, FRMS is a shared responsibility of employers and employees. Therefore, FRMS can be seen as a concrete way to engineer resilience as it requires the organisation to adjust its functioning by re-introducing safety managed by humans in addition to safety by regulations. This section presents an application of FRMS for specific cases of adaptation to the European rest requirement scheme. The whole process of the FRMS is described, from the use of predictive models of fatigue to minimise the risk at the aircrew scheduling stage to the development of fatigue-related indicators. The principles of the FRMS are discussed from the Resilience Engineering perspective.

Fatigue and Safety

Fatigue is known to be a major risk for safety in aviation. It has been classified as one of the 'most wanted' factors by the National Transportation Safety Board (NTSB) and it has been identified as the major cause of several serious incidents (NTSB, 1999) and accidents (NTSB, 1994). Although there is not a universal consensus on what fatigue is, the International Civil Aviation Organization (ICAO) gives a definition that covers most of the operational aspects that can affect aircrews: '*A physiological state of* *reduced mental or physical performance capability resulting from sleep loss or extended wakefulness and/or physical activity that can impair a crew member's alertness and ability to safely operate an aircraft or perform safety related duties.*' From this definition, it is clear that fatigue is affected by two main factors.

1. Sleep and circadian rhythms which are influenced by the hours of work.
2. Workload which is influenced by the nature of the work.

Generally, two types of fatigue are distinguished: an acute fatigue occurring when there is inadequate time to rest and recover from the work period and a chronic fatigue resulting from an insufficient recovery from acute fatigue over time. The latter is probably the least studied but may be critical because it can result from small sleep restrictions repeated over days or weeks. The cumulative effects of these small repetitive sleep losses are known to be equivalent to a total sleep restriction. People are, however, generally not aware of their negative impacts on performance and this can lead to less cautious behaviour (Van Dongen et al., 2003).

Even if the detrimental effects of fatigue on human performance have been well established in laboratory settings (Lamond and Dawson, 1999), only a few studies have investigated its effects on real situation, especially on aircrew work (Foushee et al., 1986; Thomas et al., 2006). The outcome of these studies suggest a complex and non-linear link between fatigue and safety; especially in a highly automated and team-work environment like aviation where aircrews might be able to develop strategies to mitigate the impact of fatigue (e.g., increase of cross check, automation use).

This non-linear link has been already hypothesised by Folkard and Åkerstedt (2004) from accident data and hours of work. A high level of alertness could lead to less cautious behaviour and therefore to an increased relative risk because of the operator's over confidence. A low level of alertness would be also associated to a higher risk because of a decrease of performance. Relative risk would be at the minimum for a medium level of alertness where individuals would be the most engaged in self-monitoring of their performance and more controlled processing of information. As already mentioned, fatigue awareness is probably one of the most important factors and might explain this complex link between fatigue and safety (Cabon, 2008).

Even if the underlying mechanism linking fatigue and safety is not fully understood, fatigue is widely recognised as a factor that may increase the risk of accidents. This is why fatigue has to be addressed as other risks in the management of safety and that is the objective of FRMS.

The Development of Fatigue Risk Management System

As the FRMS is intended to be an integrated part of the Safety Management System (SMS), the FRMS has been structured around the four essential components of the SMS as described by ICAO (ICAO 2008):

- safety policy and objectives
- safety risk management
- safety assurance
- safety promotion.

Figure 6.1 describes the whole process and summarises the contents of those four components in terms of FRMS.

The following sections provide more details and practical examples of how this process can be implemented by an organisation. It is the result of an on-going project (the STARE project, Cabon et al., 2008) funded by the DGAC that is intended to provide guidelines to help airlines to implement FRMS and for the regulator to oversee the FRMS implementation process. This work is performed by a consortium of experts in Human Factors, fatigue, safety and aviation with the collaboration of three regional airline partners, Airlinair, Britair and Regional.

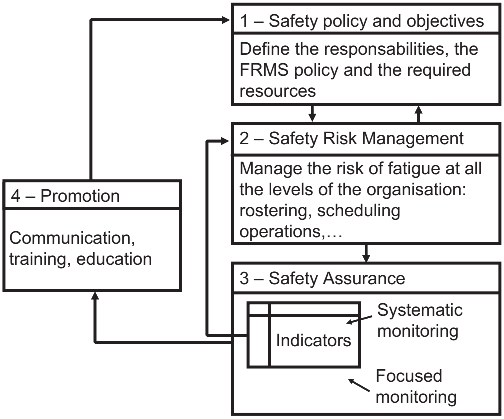
1 – Safety policy and objectives
Define the responsabilities, the FRMS policy and the required resources
2 – Safety Risk Management
Manage the risk of fatigue at all the levels of the organisation: rostering, scheduling operations,…
3 – Safety Assurance
Indicators
Systematic monitoring
Focused monitoring
4 – Promotion
Communication, training, education

Figure 6.1 The whole Fatigue Risk Management System

Safety Policy and Objectives

This first component is fundamental as it sets all the chains of responsibilities in the organisation, allocates the required resources to manage the FRMS and fosters a just reporting culture to ensure that the aircrew will feel free to report any situation where fatigue may have played a significant role.

As aircrew fatigue might be impacted by decisions taken at every level of the organisation, it is necessary that the FRMS be managed by a network within the organisation. This must include people from the top management to the operational management, including the marketing department as they are directly impacted by the design of rosters.

Another aspect that has to be covered at this level is the definition of a clear policy about issues that may impact fatigue. Examples are: the rules related to the right to refuse to report for a duty, the link between the remuneration and the most disruptive schedules, and the management of aircrew *desiderata* (e.g., to avoid aircrew to cumulate duty time to generate free time).

Fatigue Risk Management

The safety risk management component can be considered as the heart of the FRMS as it defines all the necessary steps to manage the risk of fatigue. This process covers four main steps.

1. Identification of fatigue factors.
2. Evaluation of fatigue risk.
3. Evaluation of safety risk.
4. Risk mitigation.

Identification of Fatigue Factors

A fatigue factor is defined as a factor that has an intrinsic potential of generating a fatigue risk, alone or combined with other factors. For example, reduced rest has a potential to increase fatigue as it reduces the number of hours of sleep opportunity. However, this factor does not always by itself lead to a reduced sleep and to an excessive fatigue. Therefore a systematic listing of fatigue factors has to be established. Two categories have to be considered (Table 6.1): the fatigue factors related to the hours of work (which impact sleep and circadian rhythms) and the fatigue factor related to the contextual aspects of the rosters (which impact aircrew workload). Of course, the list provided in this table is only indicative and has to be adapted by the airlines according to its own operations. The relative weight given to each item also has to be adjusted.

Table 6.1 Fatigue factors related to the hours of work

Fatigue Factors related to the hours of work	Contextual Fatigue Factors
Reduced rests Split Duties Duties starting before 06:00 Night Duties Working periods with more than 5 working days	Number of flights Complexity of the airport procedures Weather conditions Frequent changes of aircraft Technical failures Accommodation facilities (comfort, noise, temperature) Quality of the ground assistance Long commuting time Personal factors

Evaluation of Fatigue Risk

This considers the probability that an aircrew reaches an excessive level of fatigue resulting from the combination of fatigue factors. Once the factors have been identified, it is necessary to consider their possible combinations. For example: a reduced rest combined with poor accommodation and a long duty is more likely to increase the risk of fatigue than a reduced rest combined with good accommodation and a short duty. At this stage, biomathematical models of fatigue can efficiently support the Fatigue Risk evaluation. Biomathematical models are able to predict the risk of fatigue associated with specific patterns of working hours. Several software tools (for a review see Neri and Nunneley, 2004) have been developed that might be usable for the design of a new roster after the introduction of a new route or as a result of a significant change of schedule. During the current project, the duty rosters that include a reduced rest have been extracted from the planning database of the three partner airlines and evaluated through a biomathematical model, the Fatigue Risk Index (Spencer et al., 2002 that make it possible to predict the estimated risk of fatigue associated with a specific work schedule. This FRI is expressed as the probability of reaching a sleepiness level which has been validated in laboratory studies as a critical level for human performance, that is, the value of 7 on the Karolinska Sleepiness Scale (Åkerstedt and Guillberg, 1990). Figure 6.2 shows the distribution of FRI scores associated to these duties.

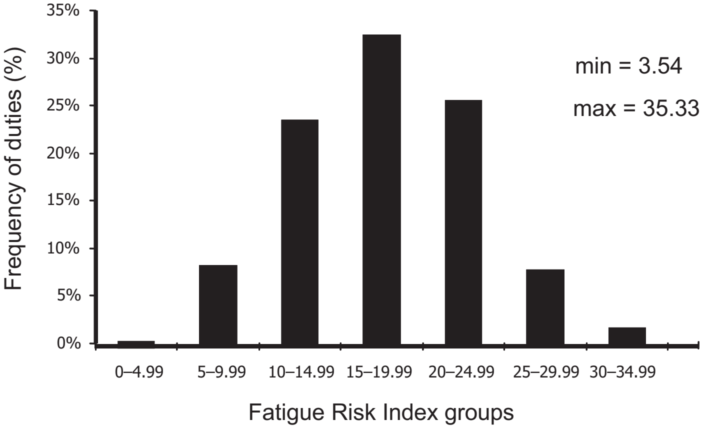

Figure 6.2 Distribution of the FRI scores associated with duties that follow a reduced rest

Note: The score is the probability of reaching a critical level of sleepiness (higher than 7 on the Karolinska Sleepiness scale)

Surprisingly, a rather large variability is observed, the scores ranging from 4.05 to 42.76. A more detailed analysis suggests that this variability is mainly due to the position of the reduced rest in the sequence of planning and therefore to a cumulative effect. For example, the fatigue risk is much lower for a reduced rest falling at the beginning of a week than for a reduced rest falling at the end of the week. This suggests that those reduced rests should be planned by taking into account the cumulative effects induced by the succession of disruptive hours of work. Interviews with the planning officers in the airlines suggest that this cumulative effect is not systematically taken into account and that the rosters are scheduled as 'isolated blocks'. However, beside the fact that the use of these models needs clear guidance, the results of their prediction should be weighted by the presence of other contextual fatigue factors. For example a low score associated with a specific hours of work could actually lead to high fatigue if the duty requires a high workload. Some of the biomathematical models make it possible to input a level of workload to impact the prediction of the model but the inputs are based on very general assumptions and are not specific to the aviation environment. This underlines the need to evaluate the fatigue risk associated with the contextual aspects of the work. This can be done by the means of questionnaires where the various constraints of the rosters (e.g., ground assistance, weather conditions, accommodations accessibility and quality) are evaluated by aircrews regarding their impact on workload and fatigue. The construction of these questionnaires requires several in-flight observations and interviews.

Evaluation of safety risk, that is, the probability that a given fatigue risk produces an undesired event (incident or accident). Taking into account the intricate links between fatigue and safety in a complex system such as aviation, this probability is a function of the nature of the operational situations. At this stage, a distinction has been made between normal, abnormal and emergency situations, the criticality of the ability of risk management being the highest in emergency situations. Therefore the objective of this step is to combine these three categories of situations (already identified by the airline) with the fatigue-risk evaluation performed into the previous step. This results in the creation of a risk matrix where the main risks are identified. For example a high risk of a Traffic Collision Avoidance System (TCAS) alert in a high traffic density combined with a high fatigue risk yield to a high safety risk. This matrix is a useful way of tracking the various risks in the operations and needs to be updated every year or after any significant change in the operations.

Risk Mitigation

The last stage of the process is devoted to the means of eliminating or reducing the risk to a level as low as reasonably practicable (ALARP). This process has to be considered at all the levels of the organisation (strategic, tactical and operational) since a decision at management level might impact fatigue at the sharp end (operators). The second aspect to consider is the risk mitigation level: suppression, reduction or prevention by means of barriers (Table 6.2). The favoured strategy is the suppression level which aims at removing the exposure of aircrew to fatigue factors (e.g., at the roster design stage). If this is not achievable (for operational or economic reasons) the reduction level has to be considered. This level aims at reducing the exposure of aircrew to fatigue factors. Finally, at the last level, when the two previous levels were not considered efficient enough to suppress or reduce the risk of fatigue, the aim is to avoid that excessive level of fatigue which produces negative effects on safety. For example, an aircrew experiencing a high level of fatigue in flight will use automation (e.g., auto-pilot) to avoid adverse effects of fatigue on flight management.

Table 6.2 Examples of actions at three decision levels and three risk mitigation levels

	Decision level		
Risk mitigation level	Strategic	Tactic	Operational
Suppression	Reduced rests or split duties as the last resort	Roster design	Refuse to report for the duty
Reduction	Remuneration policy. Training of schedulers officers	Scheduling, accommodation, *desiderata* management	Life hygiene, fatigue and sleep management
Barriers	N/A	Procedures design	Automation use, task rotation

Safety Assurance

This component is intended to continuously evaluate the efficiency of the FRMS through a monitoring of various indicators. Two categories of monitoring are proposed: a systematic and a focused monitoring.

Systematic Monitoring

Systematic monitoring includes the collection and analysis of existing safety data, essentially the Air Safety Reports and the Flight Data Monitoring.

Air Safety Reports (ASR) are short reports written by the captain to report any safety events occurring during the flight. They are mandatory and transmitted by the airline to the Civil Aviation Authority (CAA). The current structure of the ASR form does not include any information on fatigue or related issues as it is mainly a factual description report.

A first data analysis on 563 ASRs collected over a 12-month period (Figure 6.3) shows, however, that there is a significant effect of reduced rest and time on duty on the frequency of ASR. Surprisingly, there is a clear decrease of ASR for the longest time on duty after a reduced rest. This quantitative processing that already gives useful results needs to be enhanced by a more qualitative approach to further explain this trend.

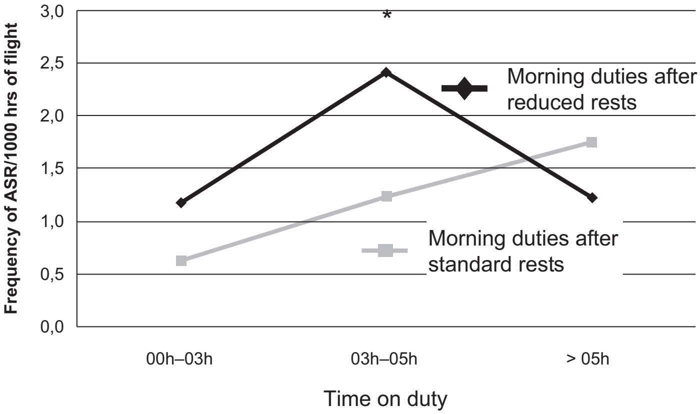

Figure 6.3 Frequency of Air Safety Reports during morning flights as a function of the nature of rest (standard or reduced) and time on duty (*n* = 230)

Therefore, it is proposed to complete this generic information with fatigue-related information. Depending on the available resources, it could be done more or less systematically. The most comprehensive means would be to add a specific 'fatigue reporting form' to the current ASR. This fatigue reporting form will collect all relevant data related to sleep quantity and quality in the last few days, fatigue evaluation at the time of the event and contextual factors that might have increased workload. Ideally, this form would be filled in systematically by both pilots. Then, only the ASR in its current form would be sent to the CAA, the ASR and the fatigue form being stored by the airline for subsequent analysis. Alternative means would be to ask aircrews to fill the form only when they estimate that fatigue was a contributing factor.

Flight Data Monitoring (FDM) is a mandatory European process for aircraft of more than 27 ton Maximum Certified Take-Off Weight that includes the acquisition, the measurement and the analysis of digital flight data (parameters, for example, speed, altitude, control wheel position, etc., but also system modes, for example, autopilot ON or OFF, etc.) in order to identify, establish probable causes for, and rectify adverse trends and deviations from accepted norms of flight operations. It is a feedback loop that provides a means for the continuous monitoring and improvement of flight operations and performance.

FDM monitors pre-determined events which are relevant for safety and performance. These events are identified by the operator regarding the aircraft and operations specifics. Events occur when one or several values exceed thresholds or when some parameters are not well-shaped at key points during the flight. To date, there is no specific event related to fatigue. STARE analysed the event occurrences and the aircrew predicted fatigue level by the means of a biomathematical model (Fatigue Risk Index), and illustrated that the occurrence of some events is significantly linked to the level of fatigue. The fatigue level where the occurrence of events is significant corresponds to only one of the two crew members being at his maximum of predicted fatigue, but not to both at the same time.

STARE results point out that FDM is a relevant tool in order to manage safety related to crew fatigue. However, a specific methodology has to be applied in order to:

- identify the aircrew fatigue level from the roster analysis;
- define the relevant severity level of events;
- classify the flights regarding the different and various fatigue factors;
- compare the occurrence of events between flights operated under 'normal' flight duty time and rest requirements and those where flexible schemes are used.

FDM gives valid indications to the operator in respect to monitoring the aircrew fatigue. The fatigue-related FDM indicators are operator-specific and defined in reference to the 'normal' rosters.

Focused Monitoring

The focused monitoring includes data that directly evaluate various aspects of the impact of work on aircrew sleep, fatigue and personal experience. This can be achieved through two main instruments: In-flight follow-up where sleep and fatigue data are collected on given roster sequences including weekly rests and survey.

- In-flight follow-up Sleep quantity and quality can be assessed through the use of a sleep log and, if possible, completed by an objective measure. Actigraphy is a light watch worn at the wrist that measures sleep quantity and quality from motor activity (Babin et al., 1997). Fatigue can be evaluated by the means of subjective scales filled in at the end of each flight. The use of validated scales like the Karolinska Sleepiness Scale (KSS) is strongly recommended. Figure 6.4 shows the mean KSS evaluation on 110 aircrews on two rosters that include reduced rests (Cabon et al., 2009). Ideally, in-flight observations are carried out to better understand the impact of contextual factors on fatigue as well as the main consequences of fatigue on aircrew activities. In the STARE project, observations were carried out over a total of 45 rosters. Those observations allowed identifying and classifying the main contextual fatigue factors. From these results, a self-assessment questionnaire was built in order to collect the relative influence of each of these factors for a large sample of aircrew. At the end of each flight, aircrews are asked to evaluate on a scale the contribution of every factor to their global fatigue.

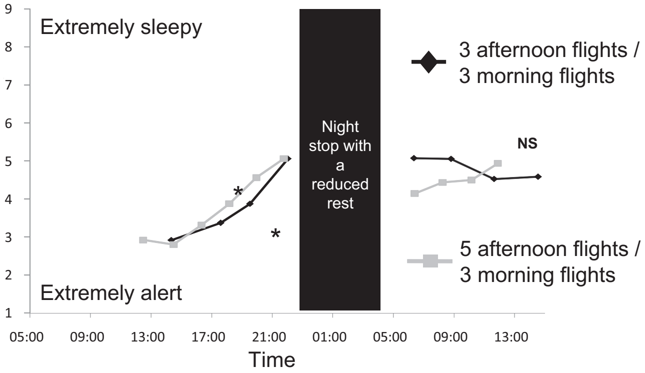

Figure 6.4 Mean KSS score on two rosters including a reduced rests. 3/3: 3 afternoon flights/night stop/3 morning flights. 5/5: 5 afternoon flights/night stop/3 morning flights ($n = 110$)

This questionnaire is currently distributed to more than 1500 aircrews. Regarding the consequences of fatigue on aircrew activities, a more qualitative approach is required as every flight is by nature specific in terms of events. Preliminary results show that some events are most frequent when aircrews feel tired. These events are, for example, omission to effectively check a check-list item, or error in Air Traffic Control (ATC) frequency reading, but also events related to the way the aircrew communicate and understand the situation. The observations show also that aircrew use behaviours and strategies (e.g., like additional cross-checks or verifications, early engagement of auto-pilot after take-off or late disengagement before landing) to manage all situations in which one or the two of the pilots have some doubts. In normal situations, mostly encountered during the study, the observed events do not impact safety, but generate more constraints and workload for the aircrew. Their potential impacts on abnormal or emergency situations cannot be directly assessed because not encountered.

- Survey Aircrews are asked to provide their own experience about fatigue such as the main causes, symptoms and coping strategies. This could represent a very useful means to track aircrew representation about fatigue and fatigue management after the implementation of the FRMS.

Monitoring Process

This section provides some suggestions on the overall monitoring process. The process makes use of all the tools discussed in the systematic and focused monitoring sections. Several scenarios can be considered or mixed to develop a coherent strategy for monitoring fatigue risk in order to manage changes and continuously improve FRMS processes.

- ‘Continuous’ mode. In this ‘basic’ mode, there is a continuous feedback from the systematic monitoring on the risk matrix. For example, a risk that was not identified in the risk matrix is added after specific events identified through the ASRs.

- 'Probe' mode. In this mode, a focused monitoring is conducted on a limited period (e.g., one month) and use to update the matrix risk.
- Proactive mode. In this mode, a focused monitoring is triggered after a significant change (e.g., introduction of a new route, schedule change). The risk matrix is updated on the basis of the results.
- Reactive mode. In this mode, a focused monitoring is triggered because of a significant change of an indicator of the systematic monitoring, e.g., frequency of ASRs and associated aircrew fatigue form increased in the last months on a specific roster.

The implementation of FRMS obviously requires a combination of tools and methods in order to manage the complexity of fatigue impact on crew safety performance.

Safety Promotion

Safety promotion has two main objectives:

- Ensure that every person in the airlines involved in the FRMS have received an appropriate training to implement and manage the FRMS.
- Ensure that every relevant person in the airlines is periodically informed of the results produced by the FRMS.

An appropriate training of the persons in charge of the FRMS is a prerequisite to the success of the FRMS. In fact, all the processes and tools that composed the FRMS need a clear understanding of what fatigue is, what the main factors impacting fatigue are, and its main consequences. Of course, depending on the involvement of a person in the FRMS, training requirements will be different. They can be divided into two levels:

- knowledge requirements;
- skills requirements.

Knowledge is basic information on a specific topic, for example, information about the sleep stages and cycles. Skills cover the ability to use techniques, tools and interpret the results, for example, the design of a sleep log and the related data processing.

Finally, the results of the FRMS should be disseminated as much as possible in the airline to provide a feedback to the aircrews. This is an important condition to maintain the necessary involvement of aircrews into the FRMS process. The feedback to the Authority also forms part of the construction of a shared understanding of fatigue management in aviation. Best practices and sharing of operational scenarios is one avenue considered to actually implement a meaningful State Safety Programme.

Conclusion

Fatigue concerns all aspects of humans at work. In such highly regulated systems as aviation, this topic perfectly illustrates the challenge to engineer 'managed' (by humans) safety in addition to 'regulated' safety in order to increase the resulting 'total' safety. The FRMS is an innovative approach to the hours of work and rest requirements focused on fatigue and safety criteria, rather than relying only on duty time regulations. Therefore, FRMS is seen as a promising way of coping with the complex management of work schedules that requires taking into account all the underlying dimensions, i.e., economic, social and safety requirements.

The on-going research has uncovered many 'no news issues' for those who question the 'Traditional Safety Perspective'. One challenge is to articulate those findings into an acceptable scheme offered by SMS, and ultimately develop pragmatic implementation guidelines for non-scientific but highly technical operational professionals. Another challenge is to develop adaptable and relevant acceptability criteria for the authority inspectors. Changing prescriptive-based for performance-based requirements involves a new way of looking at authority oversight. For example, expertise in risk management might be needed in addition to airline operations domain expertise. In addition meta-criteria should be developed to evaluate the appropriate characteristics of the airline in order to make sure that it is ‘engineered’ to be resilient. Theoretical paradigms such as complexity theories might provide clues to that extent (e.g., the DGAC’s current research study ‘PREDIR’). The adaptability and the acceptability of those guidelines will be part and parcel of the success in engineering FRMS as the support for the new resilience level for this part of the aviation business.

Those results also illustrate the various resilience levels that exist even for ultra safe industries such as aviation. In order to gain operational flexibility in response to competitive pressures in a challenging economic context for this industry, airlines are negotiating adaptation which might involuntarily decrease the resilience of front line actors. This raises the scientific question of competing or converging resilience for systems and individuals: can we achieve both or does the system resilience exclude the resilience at the individual level? The answer may be crucial for fatigue management.

Acknowledgements

This work was conducted under a grant of the Direction Générale de l’Aviation Civile with the active cooperation of the French regional Airlines: Airlinair, Britair, Regional.

Disclaimer

The ideas expressed in this paper only reflect the opinion of the authors and must not be considered as official views from any national or international authorities or official bodies to which the authors belong.

Chapter 7

Practices for Noticing and Dealing with the Critical. A Case Study from Maintenance of Power Plants

Elizabeth Lay

> If some evil genius were given the job of creating an activity guaranteed to produce an abundance of errors, he or she would probably come up with something that involved the frequent removal and replacement of large numbers of varied components, often carried out in cramped and poorly lit spaces … and usually under severe time pressure. …those who started a job need not necessarily be the ones required to finish it … a number of different groups work on the same item of equipment (Reason and Hobbs, 2003: 1).

This is an apt description of turbine maintenance work. This type of maintenance work can be fraught with rework and incidents that result from human error. The impact to power producers (utilities and independent power providers) of a maintenance service provider's failure to perform to plan during a maintenance outage can result in high losses as every day the power plant is down for maintenance, they are not selling power. For nuclear power plants, this loss can be in excess of one million US dollars in lost revenue per day. Thus, the ability of a maintenance service provider to perform to plan is critical and can be the differentiating factor in the choice of who will perform the work. This chapter covers a case study in power plant maintenance wherein principles of Resilience Engineering were used to design practices to notice risk profile changes and then move into different actions to reduce the risk in order to improve performance to plan. Implementation of the concept of 'pinging' is described. 'Pinging' is the proactive probing for risk profile changes. The steps and lessons learned for implementing 'pinging' to notice critical situations will be shared along with the design of a menu of actions to prevent the situation from turning into a 'high loss' event.

Introduction

In high risk, high pressure, complex work such as the maintenance of power plants, quality and safety incidents can occur and sometimes be extremely costly for both the service provider and the customer. Thus, performing work consistently and predictably with few incidents can be the most important differentiating factor in the choice of service provider. Reactive safety and quality programs are often limited in scope and tend to be micro-focused on specific, historical incidents or trends. Principles of Resilience Engineering can be applied to design a broad, proactive strategy for noticing the critical and moving into different actions before high loss situations occur.

Business Background

Siemens is one of the world's largest companies in the field of electrical engineering and electronics. About 400,000 employees develop and manufacture products, design and create systems and plants, and provide customized services in Industry, Energy, and Healthcare. The practices shared in this chapter were developed in the Energy Service business in the Americas where maintenance has been performed on more than 1000 turbines ranging in size from 50 to 1000 MW every year. This maintenance work is performed on nuclear and coal fuelled steam turbines and gas turbines that produce electricity.

Loss Control Philosophy

Loss is defined as the avoidable waste of *any* resource (Bird et al., 2003). Losses can result from safety and quality incidents and inefficiencies in work. The underlying management systems and possible breakdowns are mostly the same for safety, quality, and efficiency thus actions and plans to remedy potential loss situations are designed without differentiating between the three domains. In some companies, loss controls are viewed as adding cost but any loss that is controlled adds directly to profit and controlling loss can be an effective way to increase profit.

The field service group began to build the 'Story of Loss' specific to the maintenance of turbines about three years ago. The Loss Control Leadership Council was formed and included field engineers, technicians, and craft workers, about 10 people in all. After coming to a common understanding of what loss was and designing a simple method to quantify common types of loss, this council recorded and quantified the loss they saw on the jobs they were on. Other team members visited additional outages and compiled loss reports through observations and interviews with outage workers. After sampling about 40 outages over a 1 year period, the types of loss and average cost per loss-incident type were determined. The teams observed that the most common type of loss, but also the least frequently reported, was rework. Rework is defined as something done more than once or non-value added activity. The second most commonly occurring loss and also the highest cost per loss incident was waiting, waiting for tools, people, the crane, permits, decisions, parts, etc. Waiting was the highest cost type of loss event, because such loss is accrued at the burn rate (burn rate is the total cost of being on the power plant site per day including the crew's pay and expenses, tools, temporary offices, etc) if the critical path was impacted, as it is in many waiting-type loss events.

The 'Story of Loss' specific to Siemens' Americas' Power Generation Field Service organization was shared widely within the organization. This story included a grounded estimate of the total amount of loss that was being incurred annually by that organization along with types of loss that had been observed during a sampling of outages, average loss amount per incident type, how frequently these incidents were occurring, and the total impact. The simple methodology for quantifying loss that was developed included typical burn rates for field service outages, as well as building an understanding among all front line workers of what burn rate and critical path were, along with blended hourly costing rates for different roles that could be used for calculating losses due to rework.

A simple 'Story of Business': Sales − Costs = Profit was combined with the 'Story of Loss' to illustrate how reducing costs (or loss) contributes directly to profit and to bring forth that a significant amount of outages would need to be performed to regain the reported annual loss in terms of profit. Considering only the reported annual non-conformance costs (which were a fraction of estimated actual loss), the number of outages that would have to be performed to bring this back to profit given current margin rates was calculated. The number was surprising to workers and helped them see the relevance of the amount of the estimated total annual loss.

This story and philosophy were shared at all levels of the field-service organization from front-line worker to management. This foundational philosophy provided a reason for action as we moved into risk and resilience design.

Highly Resilient Organizations

Highly Resilient Organizations can be characterized by the following four behaviors.

- They *anticipate* critical disruptions and situations and their consequences.
- They *notice* the critical disruptions and situations when they occur.
- They *plan* how to respond.
- They *adapt* and move into different actions.

Design of the strategy to become more resilient included tactics in each of these domains of action.

Resilience has been defined as a measure of the ability to respond to change. Highly Resilient Organizations are able to recover rapidly when work is disrupted and are able to respond to the unexpected in a way that minimizes loss or increases gain. Being resilient includes the ability to keep working after major disruption or during continuous stress and disturbances, not merely by responding or reacting to what happens but by adjusting how work is done, or moving into different actions.

The response of an organization to stress is strikingly similar to the response of a ductile metal to stress (Figure 7.1) (Woods and Wreathall, 2008). For a ductile metal, as load (or stress) is increased, it is able to recover or return to its original form when the load is removed, up to a point. This point is the yield point. Beyond the yield point, as load is applied, the material begins to permanently deform until it reaches the point where it fractures. For an organization, as demand or load on people increases, they can respond well and handle the stress up to a certain point. Beyond this point, the risk of loss increases as people reach a limit where they are overloaded and working beyond their capacity; their mental functioning is degraded and errors are more likely. They may even reach a point where they are no longer able to cope.

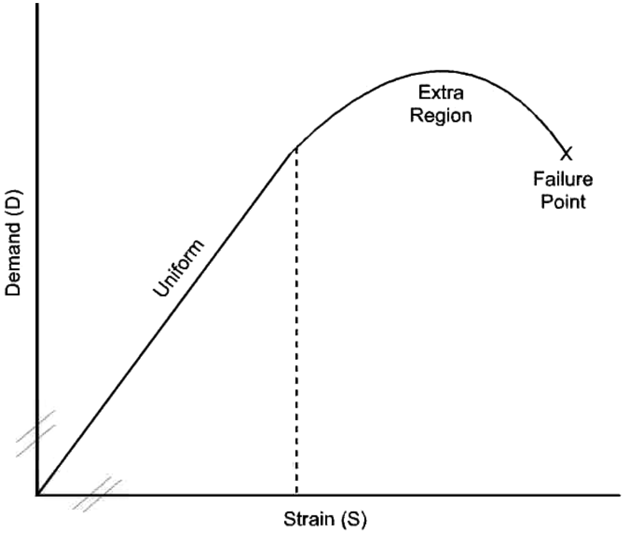

Figure 7.1 A ductile metal stress-strain curve is representative of an organization's response to stress. Highly Resilient Organizations notice when approaching the yield point or when things are taking a turn for the worse

Highly Resilient Organizations notice when people are approaching a 'yield point' and move into different actions (Figure 7.2). Building the skill and designing processes for improving 'noticing' are possible and one approach, 'pinging', is described later in this chapter. Highly Resilient Organizations move into different actions (adapt) to expand their capacity to react, extending their ability to respond to disruptions. They may remove some of the load, or stress, from the people involved in the critical situation, enabling them to return to a mode where they are able to function effectively. Some methods to do this are shared in the menu of possible solutions in the 'Adapting' section of this chapter.

Anticipate

An 'outage' involves crews of 30 to 100 or more people mobilizing to a power plant site to disassemble, inspect, and reassemble a complex machine (a turbine, for example). The work requires many specialty tools often shipped in on several tractor trailers, involves assembly of large, expensive, complex parts with very tight clearances and close tolerances, and lifting heavy components (a typical turbine rotor can weigh 50 to 80 US tons). The work is often done in extreme conditions of heat or cold, during 12 hour shifts, working 7 days a week, under extreme schedule pressure. This business has been in a growth mode with human resources being a critical factor in staffing the work that is accepted each season. It is common in this business that the teams are a mix of employees and contractors of varying levels of experience and skills, many of whom meet each other for the first time when arriving for work on the job site. Given the complexity of the work and the mix of workers, there can be many opportunities for incident likely situations to arise.

Work on the operational risk management program which includes designing practices to become more resilient began almost two years ago. The initial focus with resilience was on noticing outages where the risk profile was changing, or had changed and then moving into different actions before significant loss occurred.

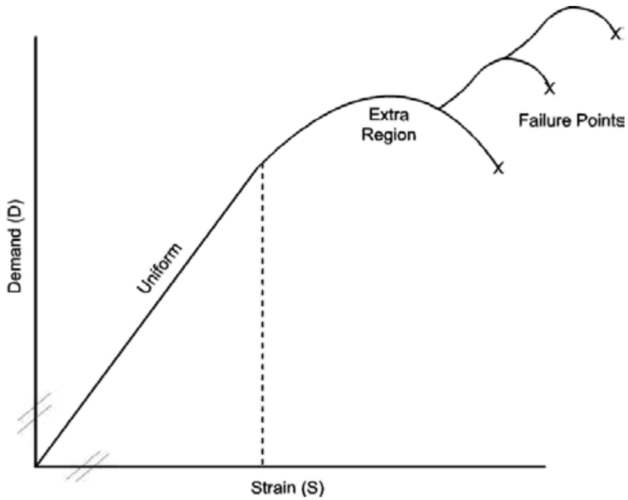

Figure 7.2 Upon noticing people are approaching a 'yield point,' Highly Resilient Organizations move into different actions extending their capacity to react

Notice

The first step on the road to become more resilient was improving 'noticing' of the critical, looking at both general situations on an outage and specific unexpected situations. One component of this was to implement a 'pinging' process. Pinging is the proactive probing for risk profile changes (Wreathall and Merritt, 2003). Through workshops with experienced project managers and operational support staff, signs that an outage could be approaching an out of control situation or risk profile change were hypothesized. Some of these potential risk factors and indicators of risk profile changes are:

- multiple issues taking crew's attention;
- progress stalling, schedule impacts, multiple delays;
- mood of project manager changes;
- specialty personnel on site longer than anticipated;
- suddenly have need for more people;
- multiple personnel being changed out;
- higher than usual amount of emergent work;

- multiple safety and quality incidents, even if minor; an increase in errors;
- common tasks not performed or performed late (such as getting permits);
- special situation with the potential to change worker's moods (working over Christmas) or risk level (Weather: storms, severe heat or cold, snow, wind, ice, hurricanes) on site;
- decline in communication, such as unreturned calls or emails;
- longer outage where potential for fatigue level is higher;
- site housekeeping has slowed or stopped.

Training was conducted with different groups of support professionals (who were not on, but were in frequent contact with, the job site). The thinking was that sometimes when you are in the heat of the battle (on the job site); it can be difficult to tell you are approaching, or are at, the 'overload point.' One of our project managers compared this to sometimes not knowing when to call the doctor when you are sick. These off-site teams were requested to refer situations where outages may be facing these types of challenges to the risk-management team. There were conversations with management on recognizing these situations. Case studies of outages that were examples of these situations were developed and shared.

Some outages were raised to the risk team's attention by operations management. None of these outages were simultaneously raised by the professionals. The potential reasons for this could have included that the professionals may not in all cases have been in a position to fully comprehend the challenge so they did not recognize the situations, or if they did notice them, they instead took immediate action and moved into a mode of helping in specific areas in alignment with their role instead of referring. The conclusion from this was that the targeted training and added pinging responsibility for specific groups of professionals did not work well in this case.

Pinging Started Narrow then Grew to a Network

The next approach was to train the entire organization from job-site clerks to front-line workers to project managers to directors on concepts of resilience, including recognizing risk profile changes, error likely situations and error likely 'climates' that could develop on specific job sites. Reason (2009: 100) noted that 'climate relates to specific workplaces and to their local management, resources and workforce. Climate is shaped by both upstream cultural factors and the local circumstances. … Unlike cultures, local climates can change quite rapidly.' Case studies based on actual incidents related to common error likely 'climates' were designed and brought forth a direct correlation between specific climates and significant loss incidents and near misses. Concluded error likely 'climates' for field service work were:

- leaders who used a 'top down' approach or intimidating style;
- leaders who were closed to listening to concerns of others on site and did not encourage questions;
- leaders who were not engaged in the work; not on the turbine deck where work was being performed;
- a site which had unclear roles / responsibilities for outage structure;
- day versus night shift competition; crew not working as a team;
- customer directing or overly involved in field service scope of work;
- leaders not familiar with current practices and cultures (contract employees);
- leaders who weren't open to help.

Noticing where these climates may exist and understanding the serious potential consequences was a first step in reducing or eliminating the potential for these climates to exist on outages. Actions the workers could take when one of these 'error likely' climates was observed were part of the training. Actions available to workers included referring the situation to management off site (anonymously if requested) and/or requesting a professional trained in risk management, human performance tools, and dealing with error likely situations to visit the site and help with the situation. It should be noted that given the severity of the potential risks associated with turbine maintenance, noticing and addressing error likely climates was not trivial.

Planning

Planning, including designing new actions, is the most challenging part of the journey to increased resilience. Once you notice that you are either in, or approaching, a difficult situation, what action do you move into? There were times when there were no extra resources available or when it was not clear what help to provide. It was determined that a variety of different solutions may be required to improve the resilience of the service business.

There is an adage that you need to accumulate and develop the power and/or knowledge and/or resources before you need them but in order to justify the cost of additional resources, you may need to prove how this 'buffering' could really change the outcome. This proof could be difficult to come by as it can be difficult to measure loss that is prevented. Siemens recently began to develop a small team, trained to act in multiple roles (field engineer, risk, safety, and quality), to pilot the buffering concept. They are being deployed to outages selectively, where the overall situation is more complex, or the complexity is increasing, and additional help may be required. If this team proves themselves to be valuable (as grounded by the number of times they are requested to help and the outcomes of outages they respond to) then the case could be built to expand the team to build additional buffering capacity, if it is needed. There have been, to date, an insufficient number of outages to which people in this role were deployed to conclusively determine how well this is working but we can say that these outages were completed without significant loss.

Adapting

It is important to note that there may not be a direct cause and effect relationship between the risk factors and subsequent risks and mitigating actions. There are limits to the capacity of any human being. As difficult situations stack-up, dealing with the situations uses some of this limited capacity and stress can reduce the capacity to act. For example, extreme weather conditions were noticed on a significant percentage of high loss outages two years in a row. It is not necessary that there be a direct correlation between the weather event and the loss incidents. Consider how working in foul weather can increase stress levels of the workers and can make work more difficult. Imagine working in frigid conditions; the worker is uncomfortable, perhaps impeded by cold weather gear, mental and physical functioning is likely reduced. Adjustments such as slowing down the work, adding stop points to warm up or hydrate, adjusting work hours, adding heaters or coolers, or adding people can help compensate. Siemens outage leadership have implemented 'stand downs,' periodic intervals during the work to check-in with the crew and have conversations between site leadership and crew for the purpose of bringing continued focus to performing safe, quality work and assessing needs of workers. Siemens also provides Human Performance Tool Kits.

Thus, when designing actions to mitigate risk profile changes, the entire situation should be considered. When assessing capacity and limits of capacity, consider how people can respond when they are close to their limits. They can become forgetful, start missing things, become more prone to outbursts of anger, show signs of fatigue and stress. Actions can be designed to add capacity (example: add workers), adjust capacity limits (example: remove stressors), or remove load, freeing up existing capacity (example: defer or move work off site).

Questions to consider for designing mitigating actions:

- What unplanned events or situations are using site leadership or worker's capacity?
- What resources or help are needed to add capacity, remove stressors, or free up existing capacity?

- Where were site personnel already with respect to their capacity limits before the situation changed? Were they already highly loaded and close to their limit?

For work during which a possible change in risk profile was observed, the first action recommended was a call to the district service manager (to whom the outage site project management ultimately reported) for their assessment of the situation. A conversation between the district manager, the regional director, and site leadership would then ensue to explore whether and what type help might be needed and where that help might be available. Key here is that a person most familiar with the situation, the customer, and the workers ultimately makes the decision on what actions to take and what help is needed. In Siemens case, this can typically be a district service manager. Siemens district offices are regionally located across the Americas and Canada. For this situation, a central support group can add value in sharing learning from outages and incidents, observing patterns and trends, and, with front-line worker's input, design actions to mitigate risk profile changes. Siemens district service managers typically have a deep and broad based background in field service work with a high level of knowledge that makes them suited to responding to difficult situations during an outage.

The following menu of possible solutions was designed by a group of experienced project managers and site supervisors:

- *Stop and assess situation*. Reassess the plan and consider where additional help is needed. Identify where the issues may be occurring.
- *Better organize site*. Evaluate parts management, tool management, roles and responsibilities, work plan/ schedule, shift turnovers, procedures, checklists, work instructions, communication plan. Pause and take the time to improve the plan.

- *Perform a Rapid Risk Assessment*. A Siemens designed process wherein the operations risk management team, site leadership, operations leadership, subject matter experts, those who may have been involved in similar situations and engineering discuss risks, mitigating actions, and clarify who the risk decision owner is.
- *Use human performance tools*. Evaluate which tools could be used that currently are not being used and whether coaching on use of tools is needed.
- *Request an Outage Specialist* to coach on human performance tools and risk and provide extra safety and quality oversight.
- *Communicate up the chain of command* to the district service office, operations management, tooling management regional directors and ask for help per the Risk Escalation policy.
- *Request commercial help*: Someone to be on site to deal with commercial situations on some complex jobs. This could give the project manager more time to work on logistics and managing the job.
- *Request logistics help*: Someone to help with parts, people, and tools, especially for emergent work or issues.
- *Begin a heightened state of coordination and help*; daily calls with those who are helping.
- *Develop resources to provide buffering*. Multi-skilled people who can travel to jobs where help is needed and act in a variety of roles.

Conclusion

Field service is already seeing benefit from applying Resilience Engineering concepts even though it is at the beginning of the journey to become more resilient. It seems that even just building the skill 'noticing' can help reduce loss. Noticing triggers action as people improvise responses even without a prescribed menu of help.

The next steps on the journey to resilience are expected to include:

- Better incorporation of practices to increase the ability to more consistently notice risk profile changes into work flow, such as adding a risk score to daily status reports.
- Continuing to improve the risk evaluations at the front of the project by integrating them with project readiness reviews.
- Continuing to observe issues that may arise and design adapting actions once critical situations are noticed.
- Considering what other types of buffering capacity could help and design the processes to bring that in.

As Siemens continues on the journey to becoming increasingly resilient, direct financial benefits, as well as enhanced customer trust and satisfaction, are expected to continue to accrue.

Chapter 8

Cognitive Strategies in Emergency and Abnormal Situations Training – Implications for Resilience in Air Traffic Control

Stathis Malakis and Tom Kontogiannis

The management of emergencies and abnormal situations in the Air Traffic Control (ATC) system entails substantial challenges for the air traffic controllers. The fundamental assumption behind the refresher training is to provide air traffic controllers with the required skills and knowledge to meet successfully a wide range of challenges imposed by emergencies and abnormal situations. Using Cognitive Systems Engineering principles, a set of cognitive and team strategies was used to explore patterns of resilience in dyadic teams of operational controllers during real and simulated emergencies in a major European Area Control Centre. The investigation of real incidents pointed to operational problems that were different from those encountered during refresher training. Results indicated a substantial gap between formal training requirements and unattained operational problems that may have safety repercussions. This chapter summarises initial findings with the potential of providing insights in cultivating sources of resilience that will supplement current refresh training.

Introduction

Annual refresher courses for operational air traffic controllers are aimed at training and equipping them with the required skills and knowledge to meet successfully the demands of emergencies and abnormal situations (EAS). There is an assumption that completing a refresher course, where classroom lessons and simulator scenarios cover standardised procedures for a range of abnormal situations, will prepare controllers to manage effectively similar situations they may encounter in actual operations.

In safety critical organisations, technological innovations are introduced for their putative quantifiable benefits of reducing workload and human error, offering a better representation of the operational environment and providing assistance to many activities of practitioners (Woods et al., 1994). In other words, technology-centred approaches offer all-encompassing solutions in a constant struggle to eradicate or contain many sources of vulnerability. Information technology enables ATC units to be equipped with high fidelity simulators that simulate almost every technical aspect of the controller's working position, including the logic of automation aids and safety nets. In this sense, simulators guide regulatory compliant training and drive an artificial need for more 'featurism' in synthetic task environments. The choice of scenarios reflects those preferred by the overarching authorities (see ICAO, 2007) whose preferences remain far from questioning.

Furthermore, studies of abnormal events in envisioned worlds discovered several hard-to-resolve decision trade-offs that call for cooperative human-system architectures (Dekker, 1996; Dekker and Woods, 1999). It is very challenging, therefore, to specify training procedures that are resilient in the face of real world ambiguities, workload demands and time constraints.

Overall, aviation refresher training appears to offer the practice of skills out of the context of real-life work while the technical and human factors aspects of flight management seem fragmented. The training emphasis on technical competencies may reflect a dominant safety paradigm in aviation that operations can be specified entirely through operating procedures that should be faithfully followed by crews. Pariès and Amalberti (2000) argued that the alternative ecological safety paradigm calls for flexible operators that can adapt procedures to situational demands; this requires a shift to new training approaches where technical skills are practised within the context of real-life demands and operators are challenged to adapt their skills, or maintain skills, in the presence of high workload and stress. The critical need for shifting to new training approaches with emphasis on non-textbook critical situations can be addressed using Resilience Engineering concepts.

Resilience represents the ability of a system to adapt or absorb disturbances, disruptions and changes and especially those that fall outside the textbook operation envelope (Woods et al., 2007). EAS represent critical situations close to the margins of safe operation that challenge the controllers' operational practices and supervisory systems. The joint human and technical system is stretched to accommodate new demands and this offers opportunities for studying aspects of system resilience. A Resilience Engineering approach should address the affordances of the system, the controller strategies and their patterns of coordination. This triad of affordances, resilience strategies and coordination has been applied in a field study of an ATC system.

An emergency presents controllers with many challenging issues. Is the situation unusual and how far to pursue monitoring of the situation? As soon as a disturbance is detected, a problem-to-be-solved is formulated and the need to re-plan for the situation becomes prominent. To respond to an emergency, controllers should demonstrate problem-detection skills and re-planning strategies. As an occurrence evolves over time, new threats may appear while the demands from current threats may change. The need for the gathering of new information to fill in the gaps, correct explanations, clarify assumptions and evaluate candidate hypotheses is amplified. This calls for strategies in recognising the situation, anticipating how the situation will evolve in future and how managing uncertainty.

On the other hand, the joint performance of controllers and supervisory systems is equally challenged in an emergency. ATC requires the synchronisation of many inter-dependent activities within a short time-window and this calls for a demonstration of joint cognitive strategies. Coordination is the main prerequisite for synchronisation but it comes at the cost of information exchange. New tasks are added and ordinary prioritisation is altered. Therefore, increased workload must be balanced by intra-team reallocation of tasks. In addition, safety critical situations are not tolerant of errors, which implies that controllers should 'engineer' their own opportunities for error detection and correction. These individual and joint cognitive strategies can be seen as important sources of resilience in the ATC system that would merit from a systematic classification.

In line with this reasoning the aim of the study presented in this chapter was twofold. First to record cognitive strategies utilised in the tasks of experienced controllers in handling of EAS. Second, to propose a Cognitive Task Analysis (CTA) method for studying how cognitive strategies emerge in a complex domain and how they can be used in the design of refresher training that provides practice conditions leading to flexible expertise.

Method

The research method was based on observations and ratings of human performance in simulator-training scenarios for operational controllers combined with participation in Team Resource Management (TRM) courses in the context of the annual refresher training. This method of identifying cognitive strategies and rating their quality was preferred to the mere analysis of incident and accident reports that focus on technical aspects and operational errors. Observational data were combined with qualitative data from briefing and debriefing simulator sessions, focused interviews with controllers and instructors, and finally an analysis of key operational documents and training curricula. These research techniques belong to the 'experiments in the field' family of methods and are based on scaled world simulations that capture critical aspects of the targeted situations (Woods and Hollnagel, 2006).

A total of 21 dyadic teams of area controllers participated in the study while attending their annual refresher training as part of their competency scheme. During the TRM course, the instructors presented a number of incidents and the controllers had vivid discussions with regard to the operational strategies adopted and the role of the context of work. A set of 14 ATC incidents were analysed involving discussions with experienced controllers and instructors. The analysis of incidents formed a unique opportunity for *listening to* how operational controllers think about them.

Our research setting was a busy European Area Control Centre where operating teams comprised two controllers. The Executive Controller (EC) was responsible for the direct control of aircraft in the sector (i.e., area of responsibility) and for carrying out the overall plan. The Coordinating Controller (CC) was responsible for establishing the overall plan for the entry and exit of aircraft in the sector and for assisting the EC in their tasks. Controllers manage air traffic in their sector by issuing arrays of complex instructions, clearances and information to flight crews. Radar is the primary sensor of the external environment (i.e., the sector and the surrounding airspace) and the main planning instrument for tactical handling of air traffic. An assortment of automated data-processing systems depicts the location of aircraft on the radar screen and enables a variety of supporting functions (e.g., Flight Data Processing).

The very essence of the area controller's task is to detect and resolve potential separation losses between aircraft in order to prevent mid-air collisions. Technically, a loss of separation is a situation where the horizontal and vertical distance between two aircraft fall below prescribed thresholds. In our research setting, the nominal separation minima between aircraft were 5 Nm horizontally and 1000 ft vertically. A loss of separation can be visualised as the overlapping of the protected airspaces of two aircraft (i.e., virtual cylinders around an aircraft with 2.5 Nm radius and 1000 ft height) while the overlap geometry determines the classification of severity.

T^{2}EAM Model

In the first stage of the study, an inventory of cognitive strategies was compiled based on a literature review from the Naturalistic Decision Making and Cognitive Systems Engineering paradigms. Four prominent sources of references were used for the individual cognitive strategies. The Recognition Primed Decision (RPD) model (Klein, 1998), the Recognition/Meta-Recognition (R/M) decision-making model (Cohen et al., 1996) and the Contingent Operator Stress Model (COSMO) decision-making model (Kontogiannis, 1999). The fourth was a model of anomaly response as a multi-threaded process (Woods and Hollnagel, 2006). These models were selected based on the importance of the cognitive strategies they integrate and the consistency of the research paradigm with the field study requirements.

For the identification of patterns of joint cognitive performance, a compilation was made of four well-established frameworks from the same research paradigms. The first one was the Anaesthetists' Non-Technical Skills, (ANTS) which is a validated and widely accepted framework (Fletcher et al., 2004). The Big Five, (Salas et al., 2005) is a teamwork model that has been developed by a critical review of empirical studies and theoretical models of teamwork, team effectiveness and team performance over the last decades. NOTECHS (Non Technical Skills) investigated possible ways to evaluate non-technical skills of multi-pilot aircrew (Flin et al., 2003).The fourth model is a taxonomy of breakdowns in shared cognition (Wilson et al., 2007).

The first stage of the study resulted in a performance model, termed as Taskwork and Teamwork strategies in Emergencies in Air Traffic Management (see T^{2}EAM in Malakis et al., 2010a, 2010b). T^{2}EAM model (Table 8.1) is an attempt for a balanced and pragmatic approach to capture resilient processes during EAS episodes in the ATC system. It is noted that the full description of the model and its internal structure is beyond the scope of this chapter.

Table 8.1 T^{2}EAM Model

Individual Strategies	Teamwork Strategies
Recognition (i.e., noticing cues, recognising states and projecting future states)	Team coordination (i.e., shared situation understanding, communication of intent, managing dependencies and avoiding information garbling)
Managing uncertainty (i.e., critiquing goals and mental models)	
Anticipation (i.e., acknowledging threats and exploiting less busy periods to plan)	Team communication (i.e., providing unsolicited information and updating situation status)
Planning for typical events and contingencies	Error management (i.e., error detection and correction)
Managing workload (i.e., prioritising tasks and coping with interruptions and distractions)	Change management (i.e., detecting and correcting problems in distribution of tasks)

For the rating of controllers' performance, we used a seven-point behaviourally-anchored scale as it was thought to give a good rating sensitivity to subject matter expert observers. The collected data were submitted to a Principal Component Analysis to establish the construct validity by revealing factor solutions that corresponded to the hypothesised models of individual and cognitive performance. The metrics of cognitive strategies were illustrated with good and poor exemplars (i.e., behavioural markers). This refinement of cognitive strategies was based on interviews with controllers and instructors so that they were able to apply this method on their own and achieve consensus in their judgement. An inter-rater validity study is currently in progress to test the screening cognitive strategy tool and promote greater use within the ATC environment. The individual and joint cognitive strategies that corresponded to T^{2}EAM are briefly analysed below.

Anticipation

Anticipation is a cognitive strategy that enables a controller to timely and accurately detect and respond to a threat. Anticipation engages with response planning during low tempo periods. It is the process of recognising and preparing for difficult challenges and brings forward the notion of threats. Threats can be defined as events or errors that occur beyond the control of the controller and must be managed in order to maintain the required margins of safety.

Recognition

Recognition is a cognitive strategy that enables a controller to timely and accurately detect early signs of an impending emergency and play out mentally the progression of events. Emergencies and abnormal situations can either occur suddenly when the flight crew formally declares an emergency or may evolve slowly over time. In the first case, recognition is effectively reduced to the elementary level of an accurate classification of the emergency type (i.e., a symptom-fault matching). In the latter case when the emergency is evolving over time, a pattern of cues should be interpreted.

Managing Uncertainty

This is a cognitive strategy that enables a controller to assemble and assess a model of the situation and establish safety related goals. Emergencies and abnormal situations are closely associated with information-based uncertainty due to their dynamics. The controller has to assemble a model of situation, formulate goals and correct any tentative explanations or assumptions, seeking information that may not be available or inaccessible. Flight crews are notoriously reluctant to provide conclusive information during emergencies and communication with the ATC is not their first priority. Even if they are willing to communicate their status, this may not be technically feasible.

Planning

Planning is a cognitive strategy that enables a controller to employ standard and/or contingency planning for the unfolding situation. Controllers have to make a plan and in certain cases to re-plan their actions in order to cope with the demands of the unfolding situation. Depending on the situation, a minimal set of prescribed action-scripts in documented forms (e.g., check-lists) are normally available in all ATC units. Controllers are trained in certain types of emergencies and this annual process is a major part of their competency scheme. Nevertheless, in many cases the need for contingency planning arises. It may be a textbook case of an abnormal situation but certain characteristics may warrant an additional form of precautionary planning in order to counteract a possible escalation of the situation.

Managing Workload

This cognitive strategy enables a controller to timely and accurately organise the required tasks and respond to interruptions and distractions. From the onset of an emergency, the workload increases significantly due to a notable change in the number of tasks, the available time and the importance of the tasks to be completed. Workload management functions as a mental task regulator enabling controllers to cope with the complexity of the situation. Issues related to switching attention between tasks and judging interruptibility are regulated by workload management.

Team Coordination

This refers to the extent to which controllers direct and coordinate other team members, establish congruence in situation assessment and clarify intent. The structure of a team (as defined by the nature of the team's tasks and their allocation) can generate lateral (intra-team) and vertical (inter-team) dependencies, which require coordination to achieve orchestrated action. The importance of coordination increases with the severity of the unfolding situation. The building blocks of coordination are the shared mental models and the intentions. The more the mental model of a team is shared, the more congruent the situation assessment of a team becomes. Communication of intentions serves to fill-in any gaps and/or tentative assumptions of the shared mental models.

Team Communication

This refers to the extent that proactive information is disseminated between controllers and regular updates are made on the situation status without disruptions and garbling. Coordination requirements generate a pressing need for communication. Communication depends mainly on information exchange and requires both sufficient time and competent cognitive resources to be accomplished. The team members exchange information to articulate their planning, their actions and their responsibilities. Therefore, the role of information exchange is crucial to the ability of team to achieve coordinated action and perform effectively in critical situations.

Error Management

This is the extent to which controllers can develop augmented monitoring strategies that enable them to detect errors and provide feedback for error correction. In handling critical situations, several errors can be committed that may complicate the situation and reduce the safety margins. Errors can be detected and corrected, not only by the individual concerned, but also more effectively through the team structure. Error detection is based on monitoring strategies that run parallel to the normal tasks, sometimes, at the cost of high cognitive resources (mainly attention and memory). In mature teams, the members employ efficient monitoring strategies that have been crafted during years of daily experience and accumulated expertise in handling system disturbances. These monitoring strategies enable them, not only to 'catch' errors promptly, but also to provide feedback for error correction without hindering the flow of tasks in the team.

Change Management

This is the extent to which controllers have developed workload-balancing strategies that enable them to detect and counteract work distribution problems. Workload is not a constant parameter and it naturally follows the changing requirements of the escalation pattern of a critical situation. The steeper the escalation pattern, the steeper the increase of the workload for the controllers. The task sequence may be altered while new tasks (those induced by the critical situation requirements) are added in the task backlog. The controllers have to manage not only the normal traffic in their sector, but also the critical situation and the interactions between them. The criticality of the situation increases and diversifies the normal distribution of work and generates imbalances between the tasks of the controllers. Therefore, a critical need arises for strategies that balance and keep the workload below saturation point for all members of the operating team.

Results

In the refresher course, controllers excelled in all cognitive strategies, especially recognition, anticipation and planning. Successful performance could be attributed mainly to the familiarisation training and the experience of controllers. Controllers were briefed about the EAS types and knew what cues to look for. They also knew how the situation could evolve over time and how to respond to the situation. The sheer amount of operational experience, combined with intensive training, resulted in excellent controller performance. High performance scores were also achieved in managing workload and uncertainty although the practised scenarios were not rated highly in terms of their demand for such strategies.

Regarding teamwork, all strategies received high scores and especially information exchange and team coordination. Expert controllers, for instance, could communicate effectively without unnecessary queries that prolonged and garbled communication. They were able to appreciate major attributes of information (i.e., criticality and time-lines) and judge the level of workload and interruptibility of other members; as a result, the EC was not distracted by the CC with non-critical information in critical phases of the situation. This success in coordination, however, provided few opportunities for practising error recovery and task change management since the controllers committed few errors.

Only two cases of separation loss were recorded in refresher training, both occurring in the 'emergency descend' scenario. An aircraft was compelled to make an emergency descend due to a technical reason and a rapid descend from the cruise level was initiated. Cues of a significant vertical deviation were evident in the radar screen. The EC failed to detect the vertical deviation in time and when it was detected, only minor heading changes to the affected aircraft were provided. These heading changes were between 5 and 15° when more than 40° would have been needed to avoid separation loss. Interestingly, the CC did not promptly inform the EC about the emergency descend and subsequently failed to question the heading changes. Conflict detection is a shared responsibility for the EC and CC positions, although the main task of conflict resolution is assigned to EC.

With regard to real incidents in the TRM course, a different picture emerged as controllers made some errors in recognising problems, anticipating threats and planning the traffic. In the real operating environment, some errors seem unavoidable especially when working under heavy workload and the influence of interruptions and distractions. It is the ability of the expert controllers to manage these errors and change their response that can contain any adverse system consequences and eventually 'engineer' resilience into the system. The result of the analysis of the TRM incidents, however, indicated that loss of separation was the result of failures in error management and task distribution. Firstly, the CC did not detect the imminent separation loss in a timely manner and did not provide feedback for error correction to the EC. Secondly, the performance of the EC in the recovery phase was less than adequate. The EC seemed to be surprised by what was happening and lost precious time in issuing questions to the pilot crews about their actions; even when the EC reacted, their interventions were not adequate to resolve quickly the situation but rather prolonged the conflict.

T^{2}EAM Model and Cognitive Task Analysis

In an attempt to operationalise the T^{2}EAM model an effort was made to develop a Cognitive Task Analysis (CTA) method by grouping cognitive strategies into three categories of performance patterns. Abstracting generic performance patterns, recurring over many variations in a particular domain, are one of the core activities of the Cognitive Systems Engineering paradigm (Woods and Hollnagel, 2006). A pattern can be described as a relational property that captures problems and opportunities arising at the intersection between sharp-end practitioners, situational demands and artefacts. The complexity of controlling a process calls for elaborate strategies of adaptation of the human-system ensemble. Three types of patterns characterise the degree of adaptation at work: Coordination (between practitioners), Resilience (between practitioners and demands) and Affordances (between practitioners and artefacts). In contrast to that, under-adaptation is characterised by their opposites: Miscoordination, Brittleness and Clumsiness. A description of generic patterns offers a good basis for the analysis of EAS demands and the analysis of cognitive strategies employed by front-end practitioners.

Patterns of Resilience

Resilience represents the ability of a system to adapt to new situations, or absorb disturbances, especially those that fall outside the usual operational envelope (Woods et al., 2007). EAS represent critical situations close to the margin of safe operation that challenge existing operational practices and supervisory systems. An emergency presents controllers with many challenging issues – is the situation unusual and how far to pursue the monitoring of the situation? As soon as a disturbance is detected, a problem-to-be-solved is formulated and the need to re-plan earlier decisions becomes prominent. To respond to an emergency, controllers should demonstrate problem-detection skills and re-planning strategies. As occurrences evolve over time, new threats may appear whilst current threats may change their demands. The need for gathering new information to fill gaps in understanding, clarify assumptions and evaluate candidate hypotheses is amplified. This calls for strategies in recognising the situation, anticipating how the situation will evolve in future, and how to manage uncertainty.

Patterns of Coordination

Coordination is based on the idea that cognition is fundamentally social and interactive (Hutchins, 1995). The concept of Joint Cognitive Systems originated from the necessity to study how work is synchronised and distributed between practitioners and artefacts, in the context of changing work demands. Coordination is an inherently complex concept that is difficult to define. ATC requires synchronisation of many interdependent activities within a short time window. However, the smooth flow of information between highly experienced practitioners can make observation of performance difficult to manage, without a deep knowledge of the domain. In the case of ATC, coordination can be exemplified by two dimensions: the internal coordination within the team and the external coordination between the teams or units. The two coordination types differ in the observability of communication and the information exchange. For example, coordination between team members is usually implicit. The CC controller understands and follows the EC by observing the radar screen and by monitoring voice communication loops between the EC and the flight crews; this could be achieved without any overt communication and at a minimum information exchange cost. On the contrary, external communication is mainly overt, as the CC has to externalise their planning and intentions to other units; this explicit form of information exchange comes at an increased cost of communication.

Patterns of Affordances

The concept of affordance is closely related to the artefacts available. In the context of Cognitive Systems Engineering, an artefact is simply something made for a specific purpose (Hollnagel and Woods, 2005). An affordance can be described as the relationship between the practitioner and the artefact. However, not all the relationships between artefacts and practitioners can be considered as affordances. In order for a relationship to be qualified as affordance, it has to provide a *fit* across the fundamental triad of practitioners, demands and artefacts (Woods and Hollnagel, 2006). An affordance is not an attribute of the artefact *per se* but an *ad hoc* support of a practitioner goal in the context of an unfolding situation. In the case of ATC, artefacts range from simple ones (e.g., using a piece of paper to note critical information, such as aircraft call-signs in a holding pattern) to complex ones (e.g., using the automated functions to amplify the detection of conflicts). In the middle of the scale, we can find the concept of an airspace sector, which provides important affordances to controllers who select and reserve a certain sector for a specific goal (e.g., for aircraft that needs to dump fuel).

The T^{2}EAM model can be used to investigate the demands of many EAS scenarios and the cognitive strategies employed by expert controllers. Challenging decisions include managing uncertainty to recognise problems, anticipating threats, standard planning, and contingency planning. Observation of expert controllers revealed that these cognitive strategies emerged in a fluid and flexible manner and shifted in response to the dynamic evolution of the scenario. The strategies are divided into three areas, reflecting adaptation or resilience, coordination and use of automated functions (affordances).

Conclusion

Several studies have suggested that the instructional facilities embedded in simulators are more important for the success of training than the simulation itself (Salas et al., 1998; Jentsch and Bowers, 1998). Instructional methods, such as training needs analysis, cognitive task analysis, scenario design, performance measurements and feedback or debriefing, are necessary to ensure mastery and evaluation of emergency response skills. Despite the earlier suggestions that aviation training should follow a systematic approach, the present study found that this systematic approach has not been applied to the refresher training in air traffic control.

Training needs analysis should be guided by CTA methods in order to specify the cognitive and team strategies that will become the focus of the training curriculum. The taxonomy of cognitive strategies specified in T^{2}EAM provide a good basis for conducting CTA and specify the cues, the challenges, the decisions, and the strategies used by experienced controllers in the course of events of a complex scenario. The proposed CTA can also help instructor to enhance the '*cognitive fidelity*' of training by identifying events in a scenario that would provide opportunities to controllers to practice specific cognitive strategies. Some strategies may occur infrequently or may be difficult to observe for an instructor (e.g., managing uncertainty) and should be provoked or triggered by specific events (e.g., withholding information from the controllers). Fowlkes et al. (1998) have developed such an Event Based Approach to Training (EBAT) that facilitates both the practice and the evaluation of cognitive strategies. This type of evaluation focuses more on the process of behaviour rather than its outcome (e.g., speed of response, number of errors). The cognitive strategies of T^{2}EAM can also provide input to the EBAT approach for the acquisition and refreshing of skills required in ATC emergencies. Hollnagel and Woods (2006) argued that we could measure the potential for resilience but not resilience itself. In line with this reasoning, we conclude that these failure-sensitive cognitive strategies provide important practical examples of the potential for resilience in two levels. First, by providing insights on how adaptations by local actors in the form of cognitive strategies are employed to support resilience in cases of safety critical events. Second by using these cognitive strategies as the foundation blocks in the development of an advanced safety training programme with the aim of cultivating sources of resilience in the ATC system.

The examination of the refresher-training course is an initial step in tailoring the training requirements to operational reality. Additional research is needed to reproduce the tentative findings presented here and derive a more informative framework of cognitive strategies. Although field studies impose several limitations related to controllability, their findings have high face validity in the practitioners. The present findings can help ATC organisations to diagnose weakness in their training and seek advice in overcoming them. The cognitive strategies can provide the foundation of an advanced safety-training course, complementary to the existing refresher course aimed at cultivating resilience. In summary, the taskwork and teamwork strategies of the T^{2}EAM model can support CTA methods in extracting situational demands and cognitive strategies that can be practiced in a refresher training course. This approach is expected to give rise to more resilient and synchronised patterns of response to EAS simulated scenarios.

This page has been left blank intentionally

PART III
Dealing with the Potential

This page has been left blank intentionally

Chapter 9
Resilience and the Ability to Anticipate

David D. Woods

Patterns of Anticipation

The ability to anticipate and adapt runs throughout the discussions of resilience in this section and throughout this book (e.g., Chapter 16). It is important at all levels of a system to act now in ways that will help maintain control despite the obstacles that it will, or can, encounter ahead. To be resilient, a system always keep an eye on whether its adaptive capacity, as it is configured and performs currently, is adequate to meet the demands it will, or could, encounter in the future. Missing or discounting signs that adaptive capacity is degrading leaves that system vulnerable to sudden collapse or failures (Woods, 2009a). The chapters in this section identify several patterns in how resilient systems may anticipate that adaptive capacity is falling, that buffers or reserves may become exhausted, that goal priorities should be changed, etc.

The first pattern is: Resilient systems are able to recognise that adaptive capacity is falling or inadequate to the contingencies and squeezes or bottlenecks ahead. For example, this ability is fundamental to avoid becoming trapped in the decompensation pattern of maladaptive behaviour noted in Chapter 10. A related example is found in studies of tipping points in naturally occurring complex systems. Scheffer et al. (2009) looked for patterns in how natural systems respond to disruptions in order to find signs that indicate when that system is exhausting its capacity to adapt and nearing collapse point. They found that a slowing time to recover from disruptions was a good indicator that the natural system in question was nearing a tipping point.

There are many facets to this ability because it concerns change over a relational property – what kinds of disruptions the system's architecture can handle (Woods and Wreathall, 2008). It is particularly important to be able to see signs of the potential for cascading effects, for example when changes create new connections and interdependencies (as covered in Chapter 13). Do new kinds or sizes of disrupting events begin to occur or begin to appear on the horizon? Is the ability to respond eroding or declining? Are perspectives overconfident and miscalibrated about the adaptive capacity of the system and how the system achieves this capacity? If so, this means the system is actually poised much more precariously than stakeholders realise. It also means the system works differently from what these stakeholders imagine and that there may be are hidden or unappreciated sources of resilience that help to preserve system function in the face of actual patterns of disturbances (Woods and Cook, 2006).

The second pattern is: *Resilient systems are able to recognise the threat of exhausting buffers or reserves*. All four chapters in this section make reference to this theme. Other studies of resilience have noted its importance as well: Cook and Rasmussen (2005) referred to the process as the risk of 'going solid'. Cook (2006) studied how intensive care units anticipated the risk of a bed crunch. Buffers can be gradually eroded over time through a series of small decisions as occurred in the events leading up to the fatal launch of the space shuttle *Columbia* (Woods, 2005). Chapter 10 examines urban fire fighting and notes that incident commanders explicitly try to avoid 'all hands' situations where they have committed all of the available resources and therefore will be unable to respond effectively to the next event or disruption. Incident commanders maintain reserves which could be deployed to fill a gap or handle a new turn of events.

From its origins, Resilience Engineering has included the problem of how to maintain buffers and reserves in the face of acute economic pressures. Interestingly, both Chapters 11 and 12 identify professionalism as a special and important resource that guards against the erosion of buffers in the face of economic and production pressures.

As Chapter 10 notes for the case of urban fire fighting, the concept of 'margins of manoeuvre' appears relevant. One can think of a facet of anticipation in light of that concept – resilient systems are able to assess how 'margins of manoeuvre' are expanding or contracting relative to the potential for surprise.

The third pattern is: *Resilient systems are able to recognise when to shift priorities across goal tradeoffs*. Studies of complex adaptive systems have revealed that tradeoffs are fundamental and inescapable – optimality-brittleness, efficiency-thoroughness, acute-chronic (Csete and Doyle, 2002; Woods, 2006a; Hollnagel, 2009). Thus, systems exist in the space defined by these trade-offs: does a system know where it is positioned in the trade-off space, can the system assess whether this position is appropriate for the context, and can the system shift its position in the trade-off space to move to better region? It is likely there are strong constraints about the structure of these tradeoff spaces and about how systems can adjust their position in these spaces based on fundamental findings about complex adaptive systems (Alderson and Doyle, 2010).

A critical indicator of resilience is how organisations manage situations where goals conflict and Chapter 11 develops and tests a way to assess how an organisation handles these kinds of situations. Chapter 12 looks at the issues of the sacrifice judgements that arise from conflicts between acute production and economic goals relative to chronic safety and equity goals. A resilient system knows when to sacrifice acute production goals and prioritise chronic safety goals. As these chapters point out, if organisations are unable to support people when they back off from economic goals in order to invest in safety (the sacrifice), the organisation will be acting riskier than it realises or wants.

The fourth pattern is: Resilient systems are able to make perspective shifts and contrast diverse perspectives that go beyond their nominal system position. Chapter 13 highlights and diagrams how this is an essential facet to judge when highly interdependent processes that cut across multiple organisations are at risk of failure. But the idea of perspective contrast is present at least implicitly in all of the chapters in this section and in just about all work on what makes systems resilient. Chapter 12 examines upward and downward interactions in resilience as an example of perspective contrast. Perspective contrast also turned out to be central to the synthesis of how adaptive systems fail in Chapter 10.

The fifth pattern is: *Resilient systems are able to navigate interdependencies across roles, activities, levels*. This facet of anticipation is succinctly conveyed in Chapter 13 at the scale of financial institutions and regulatory bodies. Interestingly, the study in Chapter 10 of a different system that operates at a different scale (urban fire fighting) also identified a need to coordinate interdependencies across roles, activities, levels in order to synchronise multiple units and to keep pace with events. Without the ability to carry out this form of anticipation, systems are at risk of the adaptive breakdown pattern of working at cross-purposes or being locally adaptive but globally maladaptive.

Both the fourth and fifth patterns in anticipation point to a new research direction that is needed to turn the ability to anticipate into control strategies. Work has begun to develop new *polycentric control architectures* that dynamically manage and adapt the relationships across diverse but interdependent roles, organisations, processes, and activities (Andersson and Ostrom, 2008; Woods, 2009b).

The final pattern is: *Resilient systems are able to recognise the need to learn new ways to adapt*. This reminds us that resilience ultimately concerns how systems learn. It is difficult to step back and reflect on how the system one is part of works in a changing, interconnected, interdependent environment, identify the weaknesses, and begin to develop new ways to work. Ultimately, as Hollnagel synthesises in the Epilogue via the Resilience Analysis Grid, resilience is about how systems learn to modulate their adaptive capacities to continuously update their fitness relative to an environment of changing pressures and opportunities.

To conclude, the chapters in this section identify several patterns about how resilient systems anticipate including:

- Resilient systems are able to recognise that adaptive capacity is falling.
- Resilient systems are able to recognise that buffers or reserves become exhausted.
- Resilient systems are able to recognise when to shift priorities across goal tradeoffs.
- Resilient systems are able to make perspective shifts and contrast diverse perspectives that go beyond their nominal position.
- Resilient systems are able to navigate changing interdependencies across roles, activities, levels, goals.
- Resilient systems are able to recognise the need to learn new ways to adapt.

One intent behind the attempts to model the dynamics of a systems' adaptive capacity is to capture general properties that can be used to understand how specific systems will behave when they encounter signs that adaptive capacity is faltering in relation to the challenges ahead (Alderson and Doyle, 2010). These models allow us to take data about what has happened in terms of data about *how* the system adapted and to *what* – and to use this to project how well operational systems are prepared in advance to handle different kinds of challenge events and surprises (Hollnagel et al., 2006).

This page has been left blank intentionally

Chapter 10
Basic Patterns in How Adaptive Systems Fail

David D. Woods and Matthieu Branlat

This chapter provides one input to resilience management strategies in the form of three basic patterns in how adaptive systems fail. The three basic patterns are (1) decompensation – when the system exhausts its capacity to adapt as disturbances/challenges cascade; (2) working at cross-purposes – when roles exhibit behaviour that is locally adaptive but globally maladaptive; and (3) getting stuck in outdated behaviours – when the system over-relies on past successes. Illustrations are drawn from urban fire-fighting and crisis management. A working organisation needs to be able to see and avoid or recognise and escape when the system is falling into one of the three basic adaptive traps. Understanding how adaptive systems can fail requires the ability to contrast diverse perspectives.

The Optimist-Pessimist Divide on Complex Adaptive Systems

Adaptive System Sciences begin with fundamental trade-offs – optimality-brittleness, (Csete and Doyle, 2002; Zhou et al., 2005) or efficiency-thoroughness (Hollnagel, 2009). As an entity, group, system or organisation attempts to improve its performance it becomes better adapted to some things, factors, events, disturbances or variations in its environment (its 'fitness' improves). However, as a consequence of improving its fitness with respect to some aspects of its environment, that entity also becomes less adapted to other events, disturbances or variations.

As a result, when those 'other' events or variations occur, the entity in question will be severely tested and may fail (this dynamic is illustrated by the story of the *Columbia* space shuttle accident, e.g., Woods, 2005.

The driving question becomes whether (and how) an entity can identify and manage its position in the trade-off space? In other words, can an organisation monitor its position and trajectory in a trade-off space and make investments to move its trajectory prior to crisis events? The pessimists on complexity and adaptive systems (e.g., Perrow, 1984) see adaptive systems as trapped in a cycle of expansion, saturation and eventual collapse. The pessimist stance answers the above questions with 'No.' Their response means that as a system adapts to meet pressures to be 'faster, better, cheaper', it will become more complex and experience the costs associated with increasing complexity with little recourse.

Resilience Engineering, on the other hand, represents the optimist stance and its agenda is to develop ways to control or manage a system's adaptive capacities based on empirical evidence. Resilience Engineering maintains that a system can manage brittleness trade-offs. To achieve such resilient control and management, a system must have the ability to reflect on how well it is adapted, what it is adapted to and what is changing in its environment. Armed with information about how the system is resilient and brittle and what trends are under way, managers can make decisions about how to invest resources in targeted ways to increase resilience (Woods, 2006a; Hollnagel, 2009).

The optimist stance assumes that an adaptive system has some ability to self-monitor its adaptive capacity (reflective adaptation) and anticipate/learn so that it can modulate its adaptive capacity to handle future situations, events, opportunities and disruptions. In other words, the optimist stance looks at human systems as able to examine, reflect, anticipate, and learn about its own adaptive capacity.

The pessimist stance, on the other hand, sees an adaptive system as an automatic built-in process that has very limited ability for learning and self-management. Systems may vary in how they adapt and how this produces emergent patterns but the ability to control these cycles is very limited. It is ironic that the pessimist stance thinks people can study and learn about human adaptive systems, but that little can be done to change/design adaptive systems because new complexities and unintended consequences will sabotage the best laid plans. Resilience Engineering admits that changing/designing adaptive systems is hard, but sees it as both necessary and possible. Resilience Engineering in practice provides guidance on how to begin doing this.

This chapter provides one input to resilience management strategies in the form of three basic patterns in how adaptive systems fail. The taxonomy continues the line of work begun by Woods and Cook (2006) who described one basic pattern in how adaptive systems behave and how they fail. The chapter also illustrates these patterns in examples drawn from urban fire-fighting and crisis management. To develop resilience management strategies, organisations need to be able to look ahead and either *see and avoid* or *recognise and escape* when they are headed for adaptive traps of one kind or another. A taxonomy of different maladaptive patterns is valuable input to develop these strategies.

Assessing Future Resilience from Studying the History of Adaptation (and Maladaptation)

The resilience/brittleness of a system captures how well it can adapt to handle events that challenge the boundary conditions for its operation. Such 'challenge' events occur (1) because plans and procedures have fundamental limits, (2) because the environment changes over time and in surprising ways and (3) because the system itself adapts around successes given changing pressures and expectations for performance. In large part, the capacity to respond to challenging events resides in the expertise, strategies, tools, and plans that people in various roles can deploy to prepare for and respond to specific classes of challenge.

Resilience, as a form of adaptive capacity, is a system's potential for adaptive action *in the future* when information varies, conditions change, or when new kinds of events occur, any of which challenge the viability of previous adaptations, models, plans, or assumptions. However, the data to measure resilience comes from observing/analysing how the system has adapted to disrupting events and changes *in the past* (Woods, 2009a: 500). Past incidents provide information about how a system was both *brittle*, by revealing how it was unable to adapt in a particular evolving situation, and *resilient*, by revealing aspects of how it routinely adapted to disruptions (Woods and Cook, 2006). Analysis of data about *how* the system adapted and to *what*, can provide a characterisation of how well operational systems are prepared in advance to handle different kinds of challenge events and surprises (Hollnagel et al., 2006).

Patterns of failure arise due to basic regularities about adaptation in complex systems. The patterns are generalisations derived from analysing cases where systems were unable to prepare for and handle new challenges. The patterns all involve dynamic interactions between the system in question and the events that occur in its environment. The patterns also involve interactions among people in different roles each trying to prepare for and handle the events that occur within the scope of their roles. The patterns apply to systems across different scales – individuals, groups, organisations.

Patterns of Maladaptation

There are three basic patterns by which adaptive systems break down, and within each, there is a variety of sub-patterns. The three basic patterns are:

- decompensation
- working at cross-purposes
- getting stuck in outdated behaviours.

Decompensation: Exhausting Capacity to Adapt as Disturbances/Challenges Cascade

In this pattern, breakdown occurs when challenges grow and cascade faster than responses can be decided upon and effectively deployed. A variety of cases from supervisory control of dynamic processes provide the archetype for the basic pattern.

Decompensation occurs in human cardiovascular physiology, for example, the Starling curve in cardiology. When physicians manage sick hearts they can miss signals that the cardiovascular system is running out of control capability and fail to intervene early enough to avoid a physiological crisis (Feltovich et al., 1989; Cook et al., 1991; Woods and Cook, 2006). Decompensation also occurs in human supervisory control of automated systems, for instance in aviation. In cases of asymmetric lift due to icing or slowly building engine trouble, automation can silently compensate but only up to a point. Flight crews may recognise and intervene only when the automation is nearly out of capacity to respond and when the disturbances have grown much more severe. At this late stage there is also a risk of a bumpy transfer of control that exacerbates the control problem. Noticing early that the automation has to work harder and harder to maintain control is essential (Norman, 1990; Woods and Sarter, 2000 provide examples from cockpit automation). Figure 10.1 illustrates the generic signature for decompensation breakdowns.

The basic decompensation pattern evolves across two phases (Figure 10.1). In the first phase, a part of the system adapts to compensate for a growing disturbance. Partially successful initially, this compensatory control masks the presence and development of the underlying disturbance. The second phase of

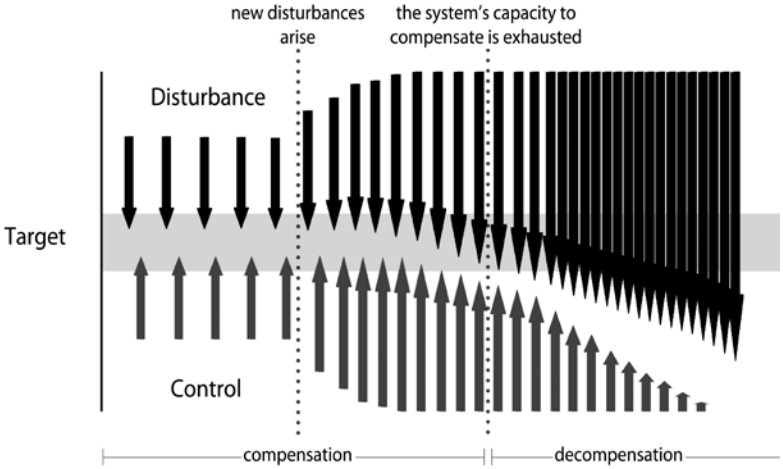

Figure 10.1 The basic decompensation signature

a decompensation event occurs because the automated response cannot compensate for the disturbance completely or indefinitely. After the response mechanism's capacity is exhausted, the controlled parameter suddenly collapses (the decompensation event that leads to the name).

The question is whether a part of the system – a supervisory controller – can detect the developing problem during the first phase of the event pattern or whether it misses the signs that the lower order or base controllers (automated loops in the typical system analysis) are working harder and harder to compensate but getting nearer to its capacity limits as the external challenge persists or grows? This requires discriminating between adaptive behaviour that is part of successful control and adaptive behaviour that is a sign of incipient failure to come.

In these situations, the critical information is not the abnormal process symptoms *per se* but the increasing force with which they must be resisted relative to the capabilities of the base control systems. For example, when a human acts as the base control system, they would as an effective team member communicate to others the fact that they need to exert unusual control effort (Norman, 1990). Such information provides a diagnostic cue for the team and is a signal that additional resources need to be injected to keep the process under control. If there is no information about how hard the base control system is working to maintain control in the face of disturbances, it is quite difficult to recognise the gravity of the situation during the phase 1 portion, and therefore to respond early enough to avoid the decompensation collapse that marks phase 2 of the event pattern. The key information is how hard control systems are working to maintain control and the trend: are control systems running out of control capability as disturbances are growing or cascading?

There are a number of variations on the decompensation pattern, notably:

- *Falling behind the tempo of operations* (e.g., the aviation expression 'falling behind the power curve;' surges in demands in emergency rooms – Wears and Woods, 2007; bed crunches in intensive care units – Cook, 2006).

- *Inability of an organisation to transition to new modes of functioning when anomalies challenge normal mechanisms or contingencies* (e.g., a hospital's ability to manage mass casualty events – see Committee on the Future of Emergency Care in the US, 2006; Woods and Wreathall, 2008 provide a general description of this risk).

Working at Cross-purposes: Behaviour that is locally Adaptive, but Globally Maladaptive

This refers to the inability to coordinate different groups at different echelons as goals conflict. As a result of miscoordination the groups work at cross-purposes. Each group works hard to achieve the local goals defined for their scope of responsibility, but these activities make it more difficult for other groups to meet the responsibilities of their roles or undermine the global or long-term goals that all groups recognise to some degree.

The archetype is the tragedy of the commons (Ostrom, 1990, 1999) which concerns shared physical resources (among the most studied examples of common pools are fisheries management and water resources for irrigation). The tragedy of the commons is a name for a baseline adaptive dynamic whereby the actors, by acting rationally in the short term to generate a return in a competitive environment, deplete or destroy the common resource on which they depend in the long run. In the usual description of the dynamic, participants are trapped in an adaptive cycle that inexorably overuses the common resource (a 'pessimist' stance on adaptive systems); thus, from a larger systems view the local actions of groups are counter-productive and lead them to destroy their livelihood or way of life in the long run.

Organisational analyses of accidents like the *Columbia* space shuttle accident see production/safety trade-offs as similar to the tragedies of the commons. Despite the organisations' attempts to design operations for high safety and the large costs of failures in money and lives, line managers under production pressures make decisions that gradually erode safety margins and thereby undermine the larger common goal of safety. In other words, safety can be thought of as an abstract common pool resource analogous to a fishery. Thus, dilemmas that arise in managing physical common pool resources are a specific example of a general type of goal conflict where different groups are differentially responsible for, and affected by, different sub-goals, even though there is one or only a couple of commonly held over-arching goals (Woods et al., 1994; Woods et al., 2010: Chapter 4). When the activities of different groups seem to advance local goals but undermine over-arching or long-term goals of the larger system that the groups belong to, the system-level pattern is maladaptive as the groups work at cross-purposes. Specific stories that capture this pattern of adaptive breakdown can be found in Brown (2005), who collected cases of safety dilemmas and sacrifice judgments in health-care situations.

There is a variety of sub-patterns to working at cross purposes. Some of these concerns vertical interactions, that is, across echelons or levels of control, such as the tragedy of the commons. Others concern horizontal interactions when many different groups need to coordinate their activities in time and space such as in disaster response and military operations. This pattern can also occur over time. A sub-pattern that includes a temporal component and is particularly important in highly coupled systems is missing the side effects of change (Woods and Hollnagel, 2006). This can occur when there is a change that disrupts plans in progress or when a new event presents new demands to be handled, among other events. Other characteristic sub-patterns are:

- *Fragmentation over roles* (stuck in silos; e.g., precursors to *Columbia* space shuttle accident, Woods, 2005).
- *Failure to resynchronise following disruptions* (Branlat et al., 2009).
- *Double binds* (Woods et al., 2010).

Getting Stuck in Outdated Behaviours: Over-relying on Past Successes

This pattern relates to breakdowns in how systems learn. What was previously adaptive can become rigid at the level of individuals, groups or organisations. These behaviours can persist even as information builds that the world is changing and that the usual behaviours and processes are not working to produce desired effects or achieve goals. One example is the description of the cycle of error as organisations become trapped in narrow interpretations of what led to an accident (Cook et al., 1998).

This pattern is also at play at more limited operational time scopes. Domains such as military operations offer a rich environment for studying the pattern. When conditions of operation change over time, tactics or strategies need to be updated in order to match new challenges or opportunities. While such decisions are made difficult by the uncertain nature of the operations' environment and of the outcome of actions, missed opportunities to re-plan constitute sources of failure (Woods and Shattuck, 2000). Mishaps in the nuclear industry have also exemplified the pattern by showing the dangers of 'rote rule following' (Woods and Shattuck, 2000). In all of these cases there was a failure to re-plan when the conditions experienced fell outside of the boundaries the system and plans were designed for. Some characteristic sub-patterns are:

- *oversimplifications* (Feltovich et al., 1997);
- *failing to revise current assessment as new evidence comes in* (Woods and Hollnagel, 2006; Rudolph, 2009);
- *failing to revise plan in progress when disruptions/opportunities arise* (Woods and Hollnagel, 2006);
- *discount discrepant evidence* (e.g., precursors to *Columbia*, Woods, 2005a);
- *literal mindedness, particularly in automation failures* (Woods and Hollnagel, 2006);
- *distancing through differencing* (Cook and Woods, 2006);
- *Cook's Cycle of Error* (Cook et al., 1998).

The three basic patterns define kinds of adaptive traps. A reflective adaptive system should be able to monitor its activities and functions relative to its changing environment and determine whether it is likely to fall into one or another of these adaptive traps. The three basic patterns can be used to understand better how various systems are vulnerable to failures, such as systems that carry out crisis management, systems that respond to anomalies in space flights and systems that provide critical care to patients in medicine. In the next section, we test the explanatory value of these three basic patterns by re-visiting a recent analysis of critical incidents (Branlat et al., 2009) that provided markers of both resilience and brittleness (Woods and Cook, 2006). Urban fire-fighting provides a rich setting to examine aspects of resilience and brittleness related to adaptation and coordination processes. Incident command especially instantiates patterns generic to adaptive systems and observed in other domains or at other scales (Bengtsson et al., 2003; Woods and Wreathall, 2008).

Illustration of the Basic Patterns

High uncertainty and potential for disruptions, new events and surprises all pose challenges for fire-fighting operations. The fire-fighting organisation needs to be able to adapt to new information (whether a challenge or opportunity) about the situation at hand and to ever-changing conditions. For example, consider the following case from the corpus (Branlat et al., 2009).

> Companies arrive on the fire scene and implement standard operating procedures for an active fire on the first floor of the building. The first ladder company initiates entry to the apartment on fire, while the second ladder gets to the second floor in order to search for potentially trapped victims (the 'floor above the fire' is an acknowledged hazardous position). In the meantime, engine companies stretch hose-lines but experience various difficulties delaying their actions, especially because they cannot achieve optimal positioning of their apparatus on a heavily trafficked street. While all units are operating, conditions are deteriorating in the absence of water being provisioned on the fire. The Incident Commander (IC) transmits an 'all hands' signal to the dispatcher, leading to the immediate assignment of additional companies. Almost simultaneously, members operating above the fire transmit an 'URGENT' message over the radio. Although the IC tries to establish communication and get more information about the difficulties encountered, he does not have uncommitted companies to assist the members. Within less than a minute, a back-draft-type explosion occurs in the on-fire apartment, engulfing the building's staircase in flames and intense heat for several seconds and erupting through the roof. As the members operating on the second floor had not been able to get access to the apartment there due to various difficulties, they lacked both a refuge area (apartment) and an egress route (staircase). The second ladder company was directly exposed to life-threatening conditions.

The three basic patterns can all be seen at work in this case.

Decompensation

The situation deteriorated without companies being able to address the problem promptly. The IC recognised and signalled an 'all hands' situation, in order to inform dispatchers that all companies were operating and to promptly request additional resources. As there were no uncommitted resources available, the fire companies were unable to respond when an unexpected event occurred (the back-draft), which created dangers and hindered the ability of others to assist. As a result, team members were exposed to dangerous conditions.

Working at Cross-purposes

Companies were pursuing their tasks and experienced various challenges without the knowledge of other companies' difficulties. Without this information, actions on the first floor worked against the actions and safety of operators on the second floor. Goal conflict arose (1) between the need to provide access to the fire and to contain it while water management was difficult, and (2) between the need to address a deteriorating situation and to rescue injured members while all operators were committed to their tasks.

Getting Stuck in Outdated Behaviour

The ladder companies continued to implement standard procedures that assumed another condition was met (water availability from the engine companies). They failed to adapt the normally relevant sequence of activities to fit the changing particulars of this situation: the first ladder company gained access to the apartment on fire; but in the absence of water, the opened door fuelled the fire and allowed flames and heat to spread to the rest of the building (exacerbating how the fire conditions were deteriorating). Similarly, the unit operating on the second floor executed its tasks normally, but the difficulty it encountered and the deteriorating situation required adaptation of normal routines to fit the changing risks.

Urban Fire-fighting and the Dynamics of Decompensation

During operations, it is especially important for the IC, constantly and correctly, to assess progress in terms of trends in whether the fire is in or out of control. To do this, the IC monitors (a) the operational environment including the evolution of the fire and the development of additional demands or threats (e.g., structural damages or trapped victims) and (b) the effort companies are exerting to try to accomplish their tasks as well as their capacity to respond to additional demands. Based on such assessments, the IC makes critical decisions related to the management of resources: redeploying companies in support of a particular task; requesting additional companies to address fire extensions or need to relieve members; requesting special units to add particular forms of expertise to handle unusual situations (e.g., presence of hazardous material).

ICs are particularly attentive to avoid risks of falling behind by exhausting the system's capacity to respond to immediate demands as well as to new demands (Branlat et al., 2009). The 'all-hands' signal is a recognition that the situation is precarious because it is stretched close to its maximum capacity and that current operations therefore are vulnerable to any additional demands that may occur. The analysis of the IC role emphasised anticipating trends or potential trends in demands relative to how well operations were able to meet those demands (see also Cook's analysis of resource crunches in intensive care units; Cook, 2006). For urban fire-fighting, given crucial time constraints, resources are likely to be available too late if they are requested only when the need is definitive. A critical task of the IC therefore corresponds to the regulation of adaptive capacity by providing 'tactical reserves' (Klaene and Sanders, 2008: 127), that is, an additional capacity promptly to adapt tactics to changing situations. Equivalent processes also play out (a) at the echelon of fire-fighters or fire teams, (b) in terms of the distributed activity (horizontal interactions) across roles at broader echelons of the emergency response system, and (c) vertically across echelons where information about difficulties at one level change decisions and responses at another echelon.

Urban Fire-fighting and Coordination over Multiple Groups and Goals

Fire-fighting exemplifies situations within which tasks and roles are highly distributed and interdependent, exposing work systems to the difficulty of maintaining synchronisation while providing flexibility to address ever-changing demands. Interdependencies also result from the fact that companies operate in a shared environment.

Several reports within the corpus described incidents where companies opened hose-lines and pushed fire and heat in the direction of others. These situations usually resulted from companies adapting their plan because of difficulties or opportunities. If the shift in activity by one group was not followed by a successful resynchronisation, it created conditions for a coordination breakdown where companies (and, importantly, the IC) temporarily lost track of each other's position and actions. In this context one group could adapt to handle the conditions they face in ways that inadvertently created or exacerbated threats for other groups. Another example in the corpus was situations where companies' capacity to fulfil their functions were impeded by actions of others. One group's actions, though locally adaptive relative to their scope, introduced new constraints which reduced another company's 'margins of manoeuvre' (Coutarel et al., 2003). This notion refers to the range of behaviours they are able to deploy in order to fulfil their functions and therefore to their capacity to adapt a course or plan of action in the face of new challenges. Such dynamics might directly compromise members' safety, for example when the constrained functions were critical to egress route management. In one case, a company vented a window adjacent to a fire escape, which had the consequence of preventing the members of another company operating on the floor above from using the fire escape as a potential egress route, should it have been needed.

Goal conflicts arise when there are trade-offs between achieving the three fundamental purposes of urban fire-fighting: saving lives, protecting property and ensuring personnel's safety. This occurs when, for example, a fire department forgoes the goal of protecting property in order to minimise risk to fire-fighters.

Incidents in the corpus vividly illustrate the trade-offs that can arise during operations and require adaptations to on-going operations. Under limited resources (time, water, operators), the need to rescue a distressed fire-fighter introduces a difficult goal conflict between rescue and fire operations. If members pursue fire operations, the victim risks life-threatening exposure to the dangerous environment. Yet by abandoning fire operations, momentarily or partially, team members risk letting the situation degrade and the situation becomes more difficult and more dangerous to address. The analysis of the corpus of cases found that adaptations in such cases were driven by local concerns, for example, when members suspended their current operations to assist rescue operations nearby. The management of goal conflicts is difficult when operations are not clearly synchronised, since decisions that are only locally adapted risk further fragmenting operations.

Urban Fire-fighting and the Risk of Getting Stuck in Outdated Behaviours

As an instance of emergency response, urban fire-fighting is characterised by the need to make decisions at a high-tempo and from a position of uncertainty. As fire-fighters discover and assess the problem to be addressed during the course of operations, re-planning is a central process. It is critical that adaptations to the plan are made when elements of the situation indicate that previous knowledge (on which on-going strategy and tactics are based) is outdated. The capacity to adapt is therefore highly dependent on the capacity correctly to assess the situation at hand throughout the operations, especially at the level of the IC. Accident cases show that the capacity of the IC efficiently to supervise operations and modify the plan in progress is severely impaired when this person only has limited information about, and understanding of, the situation at hand and the level of control on the fire.

Given the level of uncertainty, this also suggests the need for response systems to be willing to devote resources to further assess ambiguous signals, a characteristic of resilient and high-reliability organisations (Woods, 2006a; Rochlin, 1999). This is nonetheless challenging in the context of limited resources and high tempo, and given the potential cost of re-planning (risk of fragmenting operations, cost of redeploying companies, coordination costs).

At a wider temporal and organisational scale, fire departments and organisations are confronted with the need to learn from situations in order to increase or maintain operations' resilience in the face of evolving threats and demands. The reports we analysed resulted from thorough investigation processes that aimed at understanding limits in current practices and tools and represented process of learning and transformation. However, it is limiting to assume that the events that produce the worst outcomes are also the ones that will produce the most useful lessons. Instances where challenging and surprising situations are managed without leading to highly severe outcomes also reveal interesting and innovative forms of adaptations (Woods and Cook, 2006). As stated previously, many minor incidents also represent warning signals about the (in)adequacy of responses to the situations encountered. They are indicators of the system starting to stretch before it collapses in the form of a dramatic event (Woods and Wreathall, 2008). To be resilient, organisations must be willing to pursue these signals (Woods, 2009a). Unfortunately, selecting the experiences or events which will prove fruitful to investigate, and allocating the corresponding resources, is a difficult choice when it has to be made a priori (Hollnagel, 2007; Dekker, 2008: Chapter 3).

Recognising what is Maladaptive Depends on Perspective Contrasts

The chapter has presented three basic patterns in how adaptive systems fail. But it is difficult to understand how behaviours of people, groups and organisations are adapted to some factors and how those adaptations are weak or strong, well or poorly adapted. One reason for this is that what is well-adaptive, under-adaptive, or maladaptive is a matter of *perspective*. As a result, labelling a behaviour or process as maladapted is conditional on specifying a contrast across perspectives.

First, adaptive decision-making exhibits local (though bounded) rationality (regardless of scale). A human adaptive system uses its knowledge and the information available from its field of view/focus of attention to adapt its behaviour (given its scope of autonomy/authority) in pursuit of its goals. As a result, adaptive behaviour may be adequate when examined locally, even though the system can learn and change to become better adapted in the future (shifting temporal perspective).

Second, adaptive decision-making exists in a co-adaptive web where adaptive behaviour by other systems horizontally or vertically (at different echelons) influences (releases or constrains) the behaviour of the system of interest. Behaviour that is adaptive for one unit or system can produce constraints that lead to maladaptive behaviour in other systems or can combine to produce emergent behaviour that is maladaptive relative to criteria defined by a different perspective.

Working at cross-purposes happens when interdependent systems do things that are all locally adaptive (relative to the role/goals set up/pressured for each unit) but more globally maladaptive (relative to broader perspectives and goals). This can occur horizontally across units working at the same level as in urban fire-fighting (Branlat et al., 2009). It can occur upward, vertically, where local adaptation at the sharp end of a system is maladaptive when examined from a more regional perspective that encompasses higher level or total system goals. One example is *ad hoc* plan adaptation in the face of an impasse to a plan in progress; in this case the adaptation works around the impasse but fails to do so in a way that takes into account all of the relevant constraints as defined from a broader perspective on goals (Woods and Shattuck, 2000).

Working at cross-purposes can occur downward vertically too (Woods et al., 2010). Behaviour that is adaptive when considered regionally can be seen as maladaptive when examined locally as the regional actions undermine or create complexities that make it harder for the sharp end to meet the real demands of situations (for example, actions at a regional level can introduce complexities that force sharp end operations to develop work-arounds and other forms of gap-filling adaptations).

This discussion points to the finding in adaptive system science that all systems face fundamental trade-offs. In particular, becoming more optimal with respect to some aspects of the environment inevitably leads that system to be less adapted to other aspects of the environment (Doyle, 2000; Zhou et al., 2005; Woods, 2006a; Hollnagel, 2009). This leads us to a non-intuitive but fundamental conclusion that all adaptive systems simultaneously are as follows (Woods, 2009b).

- *Well-adapted* to some aspects of its environment (e.g., the fluency law—'well'-adapted' cognitive work occurs with a facility that belies the difficulty of the demands resolved and the dilemmas balanced; see Woods and Hollnagel, 2006),
- *Under-adapted* in that the system has some degree of drive to learn and improve its fitness relative to variation in its environment. This is related in both intrinsic properties of that agent or system and to the external pressures the system faces from stakeholders.
- *Maladapted* or brittle in the face of events and changes that challenge its normal function.

This basic property of adaptive systems means that linear causal analyses are inadequate for modelling and predicting the behaviour of such systems. Adaptive systems' sciences are developing the new tools needed to accurately model, explain and predict how adaptive systems will behave (e.g., Alderson and Doyle, 2010), for example, how to anticipate tipping points in complex systems (Scheffer et al., 2009).

Working organisations need to be able to see and avoid or recognise and escape when a system is moving toward one of the three basic adaptive traps. Being resilient means the organisation can monitor how it is working relative to changing demands and adapt in anticipation of crunches, just as incident command should be able to do in urban fire-fighting. Organisations can look at how they have adapted to disruptions in past situations to estimate whether their system's 'margins of manoeuvre' in the future are expanding or contracting. Resilience Engineering is beginning to provide the tools to do this even as more sophisticated general models of adaptive systems are being developed.

This page has been left blank intentionally

Chapter 11
Measuring Resilience in the Planning of Rail Engineering Work

P. Ferreira, J.R. Wilson, B. Ryan and S. Sharples

The significant pressures under which UK rail infrastructure currently operates provide ample research grounds for the field of resilience engineering. One of the areas on which these pressures mostly impact is the planning and delivery of engineering work. Resilience engineering was proposed as a framework for research aiming to improve the ability of the organisational system responsible for the planning of all engineering work to respond to these pressures. Within this scope, an approach to measuring resilience was developed by means of a questionnaire. A factor analysis method was used to identify underlying trends from the questionnaire data, which could then potentially be used as measurable aspects of resilience in rail engineering planning.

Introduction

The demand for increased capacity of the UK rail network has generated growing pressure to improve the planning and delivery of engineering work. As the owner of the UK rail infrastructure, Network Rail faces the challenge of delivering increasing volumes of work (maintenance, enhancements and renewals) within more diverse and shorter opportunities for access to the infrastructure, while maintaining the safety performance standards imposed by the regulatory bodies. A balance between productivity pressures and the assurance of the required safety standards has become critical for the sustainability of the rail organisation.

Resilience engineering was proposed as a framework for research aiming to understand and improve the planning system for rail engineering work delivery and its protection. The purpose is to assess the preparedness of the system, not only to respond to unforeseen (and unforeseeable) events, but also to manage known threats and pressures. This is to be achieved by looking at what aspects provide the planning system with potential for resilience (Woods, 2006a) as well as those that may erode this potential. More precisely, this research contemplates the following four major steps:

- The identification of key aspects of system operation by means of an interview process (Ferreira et al., 2008).
- The identification and assessment of parameters which describe planning performance based on analysis of the industry's historic data records.
- The identification and assessment of resilience factors applicable to the context of rail engineering planning.
- The comparison between different geographical areas of the railways in terms of system operation, planning performance parameters and resilience factors.

Given that the context of this book is addressing practical aspects of resilience, only the third work stream concerning the assessment of resilience factors will be discussed in the chapter.

Measuring Resilience Factors

A questionnaire was developed according to key concepts that characterise a resilient and a non-resilient system. After implementing the questionnaire at national level, a factor analysis was applied to the data, aiming to extract underlying trends as indicators for the level (and type) of resilience maintained by the system.

Questionnaire Design

The questionnaire has three sections. The first section aims at the assessment of resilience factors and will be the object of this discussion. The remaining two sections are dedicated to the assessment of the planning aspects identified and discussed throughout the interview process. These two sections will not be addressed here.

Woods (2006a) and Wreathall (2006) provide a broad range of concepts as indicators for the presence or absence of resilience in systems. Similar to the approach followed by Mendonça (2008), this was considered an obvious starting point for any attempt towards measuring resilience. Table 11.1 shows the concepts extracted from the literature sources and provide a brief description for each.

Table 11.1 Resilience concepts (from various chapters in Hollnagel et al., 2006)

Concepts	Description
Ability to adapt to changing conditions	The system has to be flexible enough to respond to external changes and pressures
Ability to cope with complexity	The system must be capable of maintaining normal operation whilst coping with changing conditions
Ability to manage continuous stresses	The system must be capable of maintaining normal operation, even when submitted to extreme pressure
Ability to respond to problems ahead of time	Preparedness - The system must be able to react before problems cause any disruption to normal operation
Learning culture	Willingness to respond to events by reforming and adapting as opposed to denying the need for change
Just culture	Support on reporting of issues throughout the organisation avoiding behaviours of culpability attribution
Ability to steer activities	The system must be able to control activities regardless of operating conditions
Appropriate level of information about performance	Awareness – The system must make available to its management appropriate levels of information regarding performance
High enough devotion to safety	Safety must be considered alongside other system goals
Buffering capacity	The system must have available the resources necessary to respond to arising problems and complex issues

While maintaining the definitions given in Table 11.1, the earlier work developed within this research, in particular the interview process (Ferreira et al., 2008), provided grounds for outlining a set of questions aimed at the context of rail engineering planning. This initial group of questions was peer reviewed by the members of the Ergonomics National Team at Network Rail in order to test their comprehensiveness as well as their meaningfulness concerning the intended concepts. This gave rise to an iterative process of revision and piloting that was concluded when the format of each question was believed to be strongly related to the underlying resilience concepts. The initial set of questions was brought down to the 22 statements shown in Table 11.2.

Questionnaire Implementation

As mentioned above, within the larger frame of this research the purpose was to compare the outcome of this questionnaire against the other work streams using a common geographical basis. To comply with this, the questionnaire was implemented at a national level, aiming to obtain a sample not only with organisational relevancy, but also with similar geographical representation as the other work streams.

Planners were asked to give their rating on a scale of 6 (1: Strongly disagree, 2: disagree, 3: Slightly disagree, 4: Slightly agree, 5: Agree, 6: Strongly agree). A sample of 105 planners was obtained from an estimated universe of 210 (due to ongoing reorganisation processes no exact numbers were available). The estimate is based on an average of 10 people at each of the 21 planning units existing at the time at national level. Of the initial 105 cases 7 were excluded from the analysis process on the basis of missing data.

Principal Components Analysis

Before undertaking any factor analysis process, basic statistics were developed in order to verify the reliability of the data as well as their suitability for factoring. Skewness and kurtosis tests were run to verify the distribution of each variable, as well as reliability tests for internal consistency of the data. In addition, the inter-item correlation matrix showed a substantial number of significant correlations and ratios for partial correlations indicated good levels for factorability of data. On the basis of this preliminary analysis, all variables were taken forward for the factor extraction.

Several factor extraction solutions and methods were explored using SPSS (Statistical Package for the Social Sciences). The selection of the most appropriate solution took into consideration the concept of 'simple structure' described by Kline (1994). The best fit to the selection criteria was a five component solution one with orthogonal rotation (Tabachnick and Fidell, 2007). Table 11.2 shows the loading factors for each variable.

Table 11.2 Matrix for extracted components

	Components				
	11	**22**	**33**	**44**	**45**
I receive feedback on the outcome of my planning	−.008	.208	**.597**	.385	−.064
I have a clear picture of how my planning contributes to the building of an integrated national delivery plan	.086	.119	.739	.064	−.106
I manage to finish whatever plans I start	.077	**.662**	−.040	.334	−.064
I have all the information I need to do my work	.026	**.827**	−.014	.130	.242
I have the information necessary to deal with unexpected situations	.255	**.760**	.257	−.055	.178
I have the information needed to detect potential planning failures	.093	**.552**	.397	−.050	.356
I have enough time to do my planning thoroughly	.197	.283	−.021	−.055	**.838**
I have enough time to reflect on my planning	.088	.278	−.021	.228	**.837**
I am encouraged to reflect on my planning	*−.039*	*−.089*	***.481***	***.453***	***.528***
I revise my planning whenever new information arises	**.510**	−.079	.158	.162	.204
I take into account a balance between safety and efficiency in my planning decisions	.199	.027	.170	**.681**	.174
I can adjust my way of working according to external pressures	.359	.177	**.494**	−.065	.232

Table 11.2 *Concluded*

	Components				
	11	**22**	**33**	**44**	**45**
I can solve problems even when pressured to deliver fast results	**.783**	.162	.044	.239	-.009
I can solve problems even when faced with unexpected situations	**.792**	.140	.004	.182	.042
I feel in control of my work activities	.429	**.641**	.011	−.118	.364
I assess the potential safety impacts for each of my planning decisions	.238	.068	−.154	**.795**	.079
I can identify when my planning decisions are pushing the boundaries of safe performance	.308	.113	−.024	**.786**	−.066
I can detect failures or errors in my planning before they create problems	**.658**	.147	.119	.384	.162
I have the support of my manager to make decisions	.234	−.100	**.655**	.004	.185
My management does not blame me for any poor outcome of my planning	.012	.046	**.720**	−.166	−.038
Because something has always gone well before, I feel confident that it will continue to go well in the future	*.363*	*.280*	*.165*	*−.238*	*−.099*
I can communicate my decisions promptly to those that rely on them	**.485**	.322	.259	.201	.094

Only loading factors above 0.400 were considered (shown in bold) and, where this led to multiple loadings, a minimum difference of 0.200 was imposed (Tabachnick and Fidell, 2007). According to these criteria the solution in Table 11.2 shows one non-loading (no loading factor above 0.400) and one cross-loading variable (more than one loading factor above 0.400 with difference between them below 0.200). The italicised characters indicate the items rejected on the basis of non-loading or cross-loading ('I am encouraged to reflect on my planning' and 'Because something has always gone well before …'). Overall, loading coefficients are significantly high, which demonstrates a strong correlation between items and their loading components.

Interpretation of the Extracted Components

The possibility of matching the initial set of questions to the 'four main resilience factors' discussed in the course of this book was considered as a starting point. The four resilience factors were defined as follows.

- Knowing what to do – The ability to respond to regular and irregular disruptions by adjusting normal functioning.
- Knowing what to look for – The ability to monitor aspects of system performance and its operating environment which are, or could become a threat in the near term.
- Knowing what to expect – The ability to anticipate developments and shifts in the operating environment on a long term basis, such as potential threats and pressures.
- Knowing what has happened – The ability to learn from experience.

The fact that the most appropriate solution was the one extracting five components, counts against a direct match to the four main resilience factors. Thus, other possible relations were investigated using literature support. Ferguson and Cox (1993) suggest two methods for naming the extracted components. Both methods resort to a sample of judges as a way to develop an independent interpretation, which makes their use time consuming and requiring a rather large number of participants.

For this research, an approach was developed on the basis of the Delphi method (Turoff and Linstone, 1975). The 25 members of the Ergonomics National Team at Network Rail were used as the 'discussion group'. Team members were asked to name each of the five groups of statements (variables loaded into each of the five components) according to what concept or idea they felt most accurately would describe that group, using as few words as possible. A total of 16 people responded with interpretations for each component. Table 11.3 summarises the expressions and concepts that were used by the majority of the respondents for each of the components.

Table 11.3 Interpretation of the extracted components

Components		Interpretation
1	I revise my planning whenever new information arises	Problem solving Flexibility Adaptability
	I can solve problems even when pressured to deliver fast results	
	I can solve problems even when faced with unexpected situations	
	I can detect failures or errors in my planning before they create problems	
	I can communicate my decisions promptly to those that rely on them	
2	I manage to finish whatever plans I start	Control Information
	I have all the information I need to do my work	
	I have the information necessary to deal with unexpected situations	
	I have the information needed to detect potential planning failures	
	I feel in control of my work activities	
3	I receive feedback on the outcome of my planning	Feedback Organisational support Role clarity and awareness
	I have a clear picture of how my planning contributes to the building of an integrated national delivery plan	
	I can adjust my way of working according to external pressures	
	I have the support of my manager to make decisions	
	My management does not blame me for any poor outcome of my planning	
4	I take into account a balance between safety and efficiency in my planning decisions	Safety Trade-offs
	I assess the potential safety impacts for each of my planning decisions	
	I can identify when my planning decisions are pushing the boundaries of safe performance	
5	I have enough time to do my planning thoroughly	Time available and management
	I have enough time to reflect on my planning	

Based on the previous feedback, the names shown in Table 11.4 were proposed by the investigator. Beyond considering the key words that were most frequently used by the respondents, whenever considered adequate, interpretations made use of concepts found in the resilience engineering literature.

Following the Delphi approach, team members were then given the opportunity to confirm or dispute the proposed names in the light of their initial interpretations. Each respondent was given a new table showing their own interpretations against the proposed names and asked whether they agree with the given name or still prefer their initial interpretation. 12 out of the initial 16 people responded to this second inquiry. Table 11.4 indicates the percentage (of the 12 respondents) of confirmations obtained for each of the proposed names.

Table 11.4 Names proposed for each component and confirmation level

	Component name	Confirmation: %
1	Adaptability and flexibility	92
2	Control	92
3	Awareness and preparedness	67
4	Trade-offs	92
5	Time management	100

The interpretation for all five components was considered valid by the majority of respondents. Nevertheless, to improve confidence in the outcome of this process, a clarification was sought whenever challenges were made.

The initial interpretations tended to favour one particular sub-group of the questions in each component. This was most evident for component 3. Comments made by respondents pointed towards the high number of questions contained in this component and the fact that these (apparently) bring together a more diverse set of issues. This accounted for a lower number of confirmations obtained for this component.

While components 1 and 2 seem to have a higher focus on personal capabilities, components 3, 4 and 5 could be seen as shifting towards a more organisational nature. The fact that the cross-loading item ('I am encouraged to reflect on my planning') refers to an organisational cultural aspect and that it loads onto components 3, 4 and 5 supports this assumption. Within this frame of mind, having the organisational support to reflect on ones planning could be an important underlying condition to allow for an adequate performance with regard to the aspects comprised in components 3, 4 and 5.

The non-loading variable ('Because something has always gone well before, I feel confident that it will continue to go well in the future') was aimed at complacency issues. As shown in Table 11.3, none of the interpretations for the extracted components allude to these issues, which would account for the non-loading of the variable.

The definitions shown in Table 11.5 aim at placing the resilience concepts found in Hollnagel et al. (2006) within the context of rail engineering planning. The earlier work streams, namely the interview process (Ferreira et al., 2008), was also used as a background for this process.

Table 11.5 Definition of components

Component name	Definition
Adaptability and flexibility	Planners are able to restructure their work (the building of a national plan for delivery) in response to pressures and adapt to new arising circumstances through problem solving
Control	People feel they have the means necessary, in particular information, to appropriately control and steer their activities
Awareness and preparedness	The system generates feedback and provides support in such a way that people have a clear view of how they should contribute towards responding to challenges
Trade-offs	Achieving a balance between safety and efficiency through decision-making. This can be interpreted in the light of the ETTO principle (Hollnagel, 2009)
Time management	Having the time to be thorough when planning decisions require it

Extracted Factors and the Potential for Resilience

The extracted components underline a relation between the questionnaire and resilience engineering constructs, hence showing a potential use as measurable factors. The integration of these factors into the original data set as new variables (using

SPSS) provides grounds for a quantified assessment. Although no reference can be given as to how much of each factor indicates a positive or negative contribute to resilience, the values obtained can be judged in the context of a geographical comparison. The outcome of the work stream focusing on planning performance will also provide ample grounds for comparison, not only on a geographical basis, but also by investigating variations in planning performance against the scores obtained on each of the extracted components.

The cross-system nature of the four resilience factors suggests that other parts of the organisation should be investigated, even if the focus is on one particular function. In other words, although this survey addresses engineering planning, it is likely that aspects such as the ability to learn ('knowing what has happened') will depend a great deal on other parts of the organisation. This can be put forward as one of the reasons why the outcome of the principal component analysis shows no direct relation to the four main resilience factors. It is likely that some of the aspects comprised within the four factors were beyond the scope of the planning system and therefore would not likely be captured by the respondents to the questionnaire.

On the other hand, the need to design questions in such a way that planners could relate them to their normal activities may be the cause of some distancing from a larger system perspective. In view of the questions format, the extracted components can only be interpreted from planners' personal perspectives, even if indicating behaviours contributing to or eroding resilience. Although requiring further investigation, the five extracted components can be viewed as aspects of planners' performance that would potentially contribute to the development of the system level abilities described by the four main factors.

Self-reporting methods, such as questionnaires, may be insufficient to provide a robust measure for resilience. However, such methods can be an efficient way of monitoring the system's behaviour in terms of resilience, particularly on a longer term basis through periodical applications. The progress of this research will determine how significant the obtained parameters are for the overall potential for resilience in rail engineering planning and delivery. This is expected to set the path for fine-tuning the questionnaire towards a more accurate measurement of resilience, not only within rail engineering but also considering its transfer to potential applications in other organisational sectors.

Acknowledgements

This project was funded and facilitated by Network Rail. The authors wish to recognise the fundamental contribution of all those who very openly cooperated with the surveys here discussed.

Chapter 12
The Art of Balance: Using Upward Resilience Traits to Deal with Conflicting Goals

Berit Tjørhom and Karina Aase

This chapter describes some of the processes involved in balancing conflicting goals (e.g., between safety and operation) in a change-intensive environment by using examples from civil aviation transport. The ability to handle multiple goals involves the use of both downward and upward resilience traits to address potential conflicts. By downward resilience, we mean that macro-level directions and solutions prepare for resilience through clear goal structures, infrastructure and procedures that handle the trade-offs between safety and efficiency. Upward resilience means that decisions made at the micro level in a system reflect a commitment to safety in case of goal conflicts. Changes, caused either by external or internal drivers, may alter these resilience traits by introducing loss of oversight. Changes made at the macro level of the system might have unintended consequences on the micro level, and vice versa. The chapter is based on studies conducted in the Norwegian civil aviation transport system.

Introduction

Even though a range of incentives exists in our society to ensure that commercial aviation operates safely (e.g., public opinion, passenger lists, lawsuits), the importance of highlighting the balance between safety and production goals is still prevalent

(Perrow, 1984). In a change-intensive environment with coexisting and conflicting pressures from macro- and micro-level actors, managers may set their priority on cost optimisation without having good aviation safety indicators to warn them about the erosion of safety margins (Rasmussen, 1997; Woods, 2005; Woods, 2006a, 2006b).

During the last decade, the civil aviation transport system has been exposed to several externally and internally motivated changes. Such changes have come in the form of new EU legislation and regulations, deregulation of the market, new business structures (e.g., mergers, restructurings, relocations) and new technologies. An increased focus on efficiency and cost reduction has been observed, leading to questions about whether this pressure has negative effects on the prioritisation of safety (Høyland and Aase, 2009; Aase et al., 2009). Historically, conflicting goals have been shown to be part of the causal explanations of several serious aviation accidents in Norway. Analysis of accident investigation reports has revealed that in the Skagerak accident (1989, 55 fatalities), pressure to uphold the flight programme due to a critical company economy was part of the accident picture. In the Namsos accident (1993; six fatalities), the investigation board recommended that the airline company's board of directors and top management clarify their principles for safety priority versus regularity, timeliness and economy (Tjørhom and Aase, 2010).

In this chapter, we want to explore how the processes of balancing conflicting goals are handled in today's aviation system and whether the balance between safety and production in a system becomes more complicated when changes, caused by either external or internal drivers, represent major elements of the context in which the system operates. The chapter uses empirical examples based on qualitative studies at different levels of the Norwegian aviation system (legislation, regulation, air traffic control, airport operation and airline maintenance). The main topics of the studies have been safety, management commitment to safety, change and safety prioritisation (Aase et al., 2009; Tjørhom and Aase, 2010, 2007; Høyland and Aase, 2009; Høyland et al., 2008; Pettersen and Aase, 2008; Pettersen, 2006, 2008; Hauland et al., 2007; Bjørnskau, 2005).

The Art of Balance

According to Reason (1990), '*All organisations have to allocate resources to two distinct goals: production and safety.*' In his opinion, these goals are agreeable in the long term, but from a short-term perspective, given a lack of resources, it appears as though production takes precedence over safety. Hollnagel describes the process of balancing safety and efficiency by using the efficiency-thoroughness trade-off (ETTO) principle (2009b, 2004, 2002). In this perspective, people adjust their work according to current conditions. It is never possible to be completely thorough, or fully efficient in view of scant resources, such as time, workforce, and money. According to Hollnagel, every work situation calls for trade-offs between thoroughness and efficiency. The trade-off tendency or favouritism of either efficiency or safety is dependent upon the dominant concern within an organisation or a system. It follows from the ETTO principle that it is never possible to maximise both efficiency and safety.

The problems with trade-offs involving safety, resilience or thoroughness are reinforced by the difficulties associated with measuring safety. Gaba (2000) points to the fact that signals of safety are weaker than signals of production and further refers to the asymmetry regarding measuring of these two goals. Where economic performance has a history of measurement for anticipation, safety often comes up short due to lack of leading indicators, thus creating difficulties in stating the relationship between resources and gains regarding safety. The picture becomes even more blurred as a result of the focus placed on 'best practice' (Hubbard, 2009) during the last decade. By collecting examples from successful organisations, one seeks to adapt to best practice standards for operation. These standards might address both safety and operation, but as Woods (2006a) pinpoints, what if there are too many goals implemented within an organisation? What if the different 'good' solutions compete and thereby create tension, or even worse, make the system less resilient?

To ensure that a system is able to handle the balance of fundamental trade-offs such as safety versus production (efficiency-thoroughness, optimality-brittleness, acute-chronic goals), one must create knowledge of the state of the art regarding safety and the ability to handle uncertainties. Where are the system's borders with regard to safety? The interactions of trade-offs create a need to consider sacrifice judgements or decisions where acute goals are sacrificed to put more emphasis/resources on achieving chronic goals like safety (Woods 2006a, 2006b). Sacrifice judgements involve the process of temporarily sacrificing acute production or efficiency related goals, or relaxing the pressure to achieve these goals, in order to reduce the risks of approaching too near safety boundaries (Woods 2006a: 32). Sacrifice judgements may occur when an approach to an airport is broken off during weather that increases the risks of wind shear, or when a take-off is delayed due to maintenance technicians' suspicion of airplane-related technical faults (Pettersen, 2008). In other contexts, safety might get sacrificed at the expense of effectiveness due to double binds created by poor accountability and brittle strategies that exacerbate goal conflicts. An aviation example is when an aircraft is de-iced and then enters the queue for take-off. The effectiveness of the de-icing agent degrades with time. Delays in the queue may raise the risk of ice accumulation. There have been several airplane crashes where, in hindsight, crews accepted delays of too great a duration and ice contributed to a failed take-off (Woods et al., 1994).

Downward and Upward Resilience

Woods (2006a) uses the phrase *cross-scale interactions* to describe the interrelations within a system. Decisions made at the strategic, or macro level of the system, might impact decisions made at the operational level or micro level, and vice versa. Woods (2006a) further operationalises cross-scale interactions by using the concepts of downward and upward resilience to describe the interrelated processes of value to resilience within a system, such as the civil aviation transport system, where decisions made at one level might have implications for system functions elsewhere in the system. Downward resilience includes macro level directions and solutions preparing for resilience through clear goal structures, infrastructure, and procedures to handle trade-offs. Upwards resilience includes decisions made at the micro level reflecting a commitment to safety in cases of goal conflicts (sacrifice judgements).

Downward resilience is of importance because the context and structures of a system either foster resilience or induce pressure towards resilient operations. For instance, will the ability of macro-level actors or 'distant supervisors' to communicate intent about goals, plans and procedures act as a downward resilience trait influencing how people at the micro level adapt to these governing tools? Local micro-level actors may use the distant macro level supervisors' or authorities' statements of intent behind goals, plans, and procedures in cases of unexpected events or changes (Shattuck and Woods, 1997; Woods and Shattuck, 2000). The absence of a clear goal structure, communication of intent behind the goals and a lack of willingness to implement adequate technology might create poor conditions for resilient operations and sacrifice judgements for front-line personnel. Safety goals should act as yardsticks meaning that deviations from the goals could appear as warning signals for operators and managers when operations exceed safety margins. Front-line operators may not be able to fully understand the consequences of a chosen deviation from prescribed rules because their actions or trade-offs are made in a specific contextual frame of reference – from their point of view in the organisation (Dekker, 2006). Repeated deviations from the prescribed design may, over time, become a new rule, which means that the design and the real operations become unequal (Snook, 2000; Vaughan, 1996, 2006). An accumulation of such deviations makes the system opaque and it becomes difficult to know if the decisions made regarding trade-offs are really sacrifice judgements

Upward resilience is of importance because local micro level actors might create resilience in a system using their experience, flexibility, and professionalism to handle the gap between rules and procedures, and the actions required to adapt to new circumstances (McDonald, 2006; Pettersen and Aase, 2008; Morel et al., 2008). These actions of the micro-level actors might be reflected in decisions made at the macro level as new strategic goals, elaboration of new procedures or implementation of new technology. The opposite, creating a threat to upward resilience is when operators and decision makers in the front line get stuck in a single problem frame and miss or misinterpret new information that should force re-evaluation and revision of the situation (Klein et al., 2005; Patterson and Woods, 2001). Research appears to indicate an emerging understanding of the gap between design, procedures and rules, and the work that is really going on in the front-line (Snook, 2000, McDonald, 2006; Pettersen, 2006; Pettersen and Aase, 2008). This gap can be described by hidden grey zones potentially inherent in design, procedures and rules that call for new ways to handle the operations (Pettersen and Aase, 2008). When situations appear that call for such flexibility, the operators and manager make sacrifice judgements (McDonald, 2006). These judgements are frequently based on experience and depend on professionalism among front-line operators. Professionalism means that within a system, there exists an ability to use experience and knowledge in addition or even instead of written procedures. In a study of professional sea-fishing skippers, Morel et al. (2008) found that they used multiple expert strategies to reduce risk without giving up on their fishing activity. They relied on a high level of adaptability, linked to an exposure to frequent and considerable risk. Such professionalism or craftsmanship serves as a buffer in situations of trade-offs between safety and production goals (Høyland and Aase 2009; Morel et al., 2008).

As we have seen, the ability to balance multiple goals involves using both upward and downward resilience traits and most importantly the interactions between them and across system levels. Changes caused by either external or internal drivers may alter these resilience traits by introducing loss of oversight or emerging risks. Let us now turn to some examples from real practice.

Traces of Balancing Within the Norwegian Aviation Transport System

During the last decade, the Norwegian aviation transport system has been influenced by numerous and extensive changes. These changes, along with the interconnectedness of the transport system, might impact the ability to handle multiple goals within the system. The complexity within the system is greater than ever, meaning that the risks associated with these interdependencies might be extensive and it could be useful to discuss them using the concepts of downward and upward resilience.

Methodology

To illustrate the issues of balancing, we will in the following use examples from different research studies undertaken within Norwegian civil aviation (Aase et al., 2009; Tjørhom and Aase, 2009, 2007; Høyland and Aase, 2009; Høyland et al., 2008; Pettersen and Aase, 2008; Pettersen, 2008; Hauland et al., 2007; Pettersen, 2006; Bjørnskau, 2005). The studies cover empirical data (collected over a period of four years, 2004–2007) from three cases that represent different levels of the aviation system.

- *The legislation/regulation case* consists of 26 interviews with inspectors, advisers and managers in the Norwegian Civil Aviation Authority (NCAA) and 12 interviews with employees in the Ministry of Transport and Communications. The objective of the study was to describe safety policies, perceptions of safety, safety practices and changes.
- *The air traffic control (ATC)/airport operation case* contains a study of five airports with 126 informants (interviews), aimed at diagnosing the safety culture as a means for improvement. The case also includes qualitative free text data concerning changes and safety aspects from a questionnaire survey, with 231 respondents (managers, planners, engineers, air traffic controllers) from ATC and airport operation.
- *The maintenance case* was carried out as an exploratory study of a line maintenance department, with participant observation, 15 interviews and a number of informal discussions. The goal was to gain insight into how safety is created and maintained through work practices at an individual/group level. The case also includes free text data from a questionnaire survey, with 283 respondents within maintenance (managers, planners, engineers, aviation technicians).

Using data from the different case studies in this study was done by searching the empirical material and the previous research publications for issues covering the topic of goal conflicts and for empirical examples on processes of balancing safety and production.

Downward Resilience?

There has been a transition towards deregulation of the aviation transport market, which has influenced the economic situation within the aviation business. Due to economic pressure, the structure of the companies has changed. Companies have been downsized, bought, and merged. Further, within the safety regulation framework, there has been a transition from national regulation towards a standardised EU framework for safety rules. The Norwegian Ministry of Transport and Communications states: '*The Ministry is responsible for the framework conditions within aviation transport in Norway.*' This general and overall statement is handled by the Ministry's subordinate agency, the Norwegian Civil Aviation Authority (NCAA), which is assigned responsibility for ensuring that civil aviation in Norway is operated safely and efficiently. This responsibility is made explicit in the NCAA's vision: '*NCAA should be an active initiator for safe and community-serving aviation services.*'

The background for this vision influenced by the current Ministry is the knowledge that civil aviation plays a more important role in the transport pattern in Norway than it does in most European countries and that civil aviation makes an important contribution to maintaining settlement and employment throughout Norway. The result is a network of 46 state airports across the entire country (approximately 4.8 million inhabitants). The objective of being both 'safe' and 'community-serving' seems to contain potential goal conflicts according to an informant from NCAA:

> Our goal is to be both community serving and contribute to an increased safety level within aviation. I do not agree with such double-edged goal, in my opinion our job should be to say NO to anything that may harm safety! But there are lots of difficult decisions regarding exemptions from rules and regulations that we have to deal with.

The Norwegian Aviation Act of 1993 is a so-called delegation act, which leaves to other institutional bodies the responsibility to elaborate on the details found in the body of rules. There are no clear statements from the Ministry on how potential conflicts should be handled by the NCAA in their activities related to supervising and ensuring compliance with regulations and conditions. The challenge of conflicting goals inherent in the statements of the Ministry and the NCAA might be even more prevalent because the political system in Norway is uncertain, consisting of many small political parties that form coalitions. In practice, this uncertain nature means a new political environment emerges every fourth year. Employees in the Ministry express the following about changes in government:

> 'The department changes colour, quite a lot of the attitudes change. But from day to day, the jobs we do are the same.'
>
> 'New government? Then we have to fling ourselves into the new government's declaration.'
>
> 'It often happens during preparation of different cases or elucidations that we become aware of the fact that what we prepare is against political decisions. It is important for us to act tidy on these issues.'

Changes in political climate might generate a change in the goals and statements of the Ministry of Transport and Communications, and consequently, the NCAA, that is, if a political party is especially focused on regional policy, the implication could be that when this party assumes power in the government, it might abandon existing plans for closing down some of the short take-off and landing (STOL) airports in Norway that do not satisfy the demands for airports within the EU or follow the international body of rules.

Because the NCAA still holds the technical competence to license the operation of STOL airports, it becomes their task to decide whether exemptions from existing regulations must be granted in order to operate these STOL airports. In the absence of an overall defined trade-off, the decision to exempt or not becomes a struggle between professional considerations and current political composition. This struggle indicates a vulnerability regarding the commitment to safety (or not) among the employees in the NCAA. It also returns to the role of the Ministry of Transportation to demonstrate a commitment to safety. As one NCAA employee said: '*We are the government's instrument for both a safe and community-serving aviation. Viewing resilience as an interrelated system, it becomes important to know the Ministry's opinion about commitment to safety.*'

This lack of clear guidance from the Ministry on how to prioritise conflicting goals is what Grote (2004, 2008) denotes as a deficiency in rules management. She describes 'rules management as a source for loose coupling in high-risk systems' (2008: 91). Rules can function as glue within organisations, which makes the working operations consistent even when workers must adapt to unfamiliar events. If rules should be resources rather than determinants for action, we must distinguish between different specification levels of rules. We can differentiate rules for goals, processes, and actions (Hale and Swuste, 1998). These three types of rules could be viewed as following an axis, where goal rules are the most strategic of the three and action rules are the most detailed.

The lack of distinct goal rules worked out by the Ministry of Transport and Communications has created an inherent tension between the double-edged objective of both 'safe' and 'community-serving'. None of the stated visions by the Ministry or the NCAA can serve as goal rules that give the organisation a common direction for making trade-offs between safety and production. Indeed, the visions of both organisations lack the dimension of giving direction for determining trade-offs between safety and efficiency. Decision makers then lack the directions that give them the power to make sacrifice judgements (Woods, 2006a). Without any clear or well-defined overall goal rules for safety from the macro level of the Norwegian aviation system, it is difficult to claim that the system has an inherent downward resilience.

Upward Resilience?

Within ATC/airport and operation, goal conflicts have been identified as being related to prioritisation between efficient traffic handling and safety. Differences between airports exist, in which some handle the possible conflicts by choosing safe work practices, some by addressing the conflict upwards in the hierarchical system, and some by providing the necessary resources for safe operations. The experience of other airports indicates that efficient traffic handling gets prioritised over safe operations, thus resulting in procedure violations (Høyland and Aase, 2008; Høyland et al., 2008). The examples show that when the operators experience a commitment to safety by their managers, they dare to make sacrifice decisions, as they do at the airports where they feel that commitment. The opposite is true at airports where the operators experience a lack of commitment from their management, and thereby tend to give efficiency precedence over safety. At the airports where economic pressure gets precedence towards safety, the employees expressed the situation as:

> 'It is not possible to get support for safety by the managers.'
>
> 'We feel pressure towards too much overtime work.'
>
> 'Operative personnel might lack of time to [resolve] safety issues caused by continual pressure towards administrative work task[s].'

Within aviation maintenance, the technicians report that formal descriptions of work are part of their knowledge base. In addition to the written procedures, they must elaborate on their problem-solving procedures. These procedures are used when situations call for flexibility. The standard operating procedures are static tools that need to be justified to keep the system resilient. Such problem-solving procedures are 'embedded in the heads and hands of the practitioners' (Pettersen, 2006, 2008). The technicians report about intuitive feelings that guide their judgements, based on years of experience, which offer them a comprehensive view of their part of production in an appropriately safe manner. According to the technicians, their freedom to choose safety over efficiency has changed. They experience conflicting goals related to keeping the aircraft safe from technical faults while simultaneously getting the aircraft operational within the time limit of its planned schedule. They report that when they experience conflicts regarding making (in their view) good trade-offs, they often resolve those conflicts by creating time spaces ('delays due to technical reasons') to ensure the airplane becomes technically airworthy (Pettersen and Aase, 2008). In the trade-off between punctuality and safety, the operating technicians were committed to making sacrifice judgements.

Due to the current change intensity of the Norwegian civil aviation system, many technicians have experienced increased demands for productivity. When they were asked about their perception of how the current changes affect safety, the following statements were frequent:

> 'Economy gets precedence over safety.'
>
> 'There is an odd mixture of safety and profit.'
>
> 'Generally increased demands for improved efficiency.'
>
> 'The trust in central management is considerably weakened caused by their one-sided focus on economy.'

Their perception of an increased focus on production is a challenge when it comes to their commitment to safe work practices. According to Woods (2006a), the frames for making sacrifice judgements have then been altered. Lacking a framework for sacrificing judgements based on clear and common goal rules that create downward resilience, the technicians must make their own action rules (Grote, 2008), as exemplified by 'delays due to technical reasons' rooted in their technical competence.

Conclusion

In the Norwegian aviation transport system, different studies have shown that there is a lack of commitment to downward resilience at the macro level, due primarily to the tension inherent in the double-edged objective of being both safe and community-serving. The prioritisation of regional policy (community-serving) and an unwillingness to develop distinct goal rules for balancing safe and community-serving air transport, place downwards pressure on the aviation system. Despite deficiencies in the downward resilience, upward resilience traits at the micro level of the aviation system seem to counterbalance the picture by characteristics such as a clear commitment to safety, sacrificing decisions, and establishing resource buffers to handle safety in critical situations. The critical issue regarding resilience in the Norwegian aviation transport system seems to be the awareness towards vulnerability caused by the system's dependency on upward resilience.

These findings have implications for different levels of the aviation transport system. We propose following actions to strengthen downward resilience.

- Development of clear safety goal rules at the governmental level.
- Downward resilience is threatened by the unwillingness to state clear goal rules at the strategic level. After years of changes within the aviation transport system, employees need clear statements that give them a framework to remain flexible and committed to safety despite economic pressure.
- The goal rules should be based on worst-case scenarios using input from the entire aviation transport system. The institutional level of the system must be responsible for collecting information regarding trends that threaten resilience.
- Development of guidelines and requirements for addressing cross-scale interactions.
- The training tools should include participants from different levels and professions.

We propose following to strengthen upward resilience.

- Foster perpetual awareness among operators.
- Without a constant unease about the way to handle an operation, one might become lost in routine and fail to notice variations. Even a seemingly insignificant variance in operation must be taken as a potential leading indicator regarding threats against resilience.
- Extend operators' collaboration with other parts of the system.

- A strong focus on professional values might have some downsides (McDonald, 2006). Within a profession, self-confidence may evolve to the level of overconfidence. In a trade-off situation, this may result in over-reliance on the individual's judgement – at the expense of cautious prudence. Technicians and airport operators might rely too heavily on experience and knowledge, thus taking unnecessary chances without fully embracing the body of rules. Interrelations necessitate an exchange of knowledge across professions.

The tension between downward and upward resilience in the aviation system that we have studied is balanced by a strong professionalism throughout the system, which functions as a buffer and makes safety goals prevalent over production goals. To uphold this art of balancing, it is in our opinion crucial to develop strong but flexible goal rules at the macro level to demonstrate a commitment to safety that micro level actors find trustworthy. At points of intensified production pressure and higher organisational tempo, extra investments in sources of resilience are required to keep production/safety trade-offs from sliding out-of-balance. In other words, safety investments are most important when least affordable (Woods, 2006b).

Chapter 13
The Importance of Functional Interdependencies in Financial Services Systems

Gunilla A Sundström and Erik Hollnagel

The events of 2007–2009 in the global financial markets clearly illustrated the need of an improved understanding of how the global Financial Services System (FSS) functions. In particular, the crisis made it clear that national FSSs, or components of such systems such as individual banks, were highly dependent on the normal functioning of other components of the global FSS. The primary goal of this chapter is to introduce a functional framework that enables a proactive identification of risk associated with outcomes of actions – either planned or already taken. Key concepts from Resilience Engineering and functional modelling are leveraged to define the approach. The primary goal of the proposed framework is to identify key functional dependencies between an individual firm's business functions and the functions that drive key behaviours of global financial markets. The rapid demise of the UK-based residential mortgage firm Northern Rock is used to illustrate the proposed framework.

Introduction

The events of 2007–2009 sent shockwaves through the global financial services industry. The shockwaves were triggered by an unexampled event (Westrum, 2006), namely the global credit market crunch and its impact on other parts of the global FSS.

Many reputable global financial services organisations (e.g., Bank of America, UBS, Citigroup and the Royal Bank of Scotland) experienced unprecedented losses and found themselves in need of financial support from national governments. Several other financial services firms suffered catastrophic business failures, for example, the US-based global investment banking firm Bear Stearns and the UK-based residential mortgage lending firm Northern Rock.

The 2007 global credit market crunch required extreme imagination to be comprehended and consequently pushed organisations outside of their experience envelope. What happened was that an event in the US sub-prime market rapidly propagated through the global FSS with an unprecedented impact on it, and eventually, on the world economy. The turmoil showed that using historical data to predict the future obviously did not provide the required forward-looking assessment of financial market behaviour (Bernstein, 2007). Unfortunately, many traditional risk metrics and forecasting techniques of the financial services industry do rely on historical data, including the Value-at-Risk (VaR) approaches used to determine market risk (Manganelli and Engle, 2001).

In this chapter we demonstrate how Resilience Engineering can provide the financial services industry with a better way to understand the potential impact of both past and future actions by identifying how system components are interconnected. The term 'functional' is used in the present context to emphasise that the focus is on capturing the behaviour of financial services functions and not on describing implementation details, that is, the mechanisms of any specific function. (cf., Merton, 1995; Merton and Bodie, 2005). Throughout the chapter, the term 'risk' will be used to denote a state in which a FSS in part, or in whole, is exposed to uncertain outcomes that can be either positive or negative.

The Financial Services System 2007–2009

The turmoil in the global financial markets has been the subject of several financial stability reports, published by central banks since 1996–1997, as well as reports published by the Financial

Stability Board (FSB), previously known as the Financial Stability Forum. A recent FSB report (2009), highlighted procyclicality as one of the major contributors to the disruptions in the global financial markets. The report defined 'procyclicality' as 'the dynamic interactions (positive feedback mechanisms) between the financial and the real sectors of the economy. These mutually reinforcing interactions tend to amplify business cycle fluctuations and cause or exacerbate financial instability' (p. 10). In other words, there are functional dependencies between the economy and the global financial markets that can provide mutually amplifying reinforcements. Thus, if global financial markets contract, the economies also tend to contract. In their report FSB states that 'amplifying feedback mechanisms can be as potent in the expansion phase of the [business] cycle as they are in downturns.'

The FSB highlighted two primary sources of procyclicality:

- Risk Management limitations: 'Measures of risk often spike once tensions arise, but may be quite low even as vulnerabilities and risk build up during the [business] expansion phase.'
- Inherent conflict between providers and users (i.e., lenders and borrowers) of funds, that is, the so-called 'Principal-Agent problem'. FSB sees this conflict as particularly difficult if there is a close link between asset valuations and funding. An example of such a conflict between lenders and borrowers is the need for individual lenders to retain capital; as a result they cannot make economic resources available to borrowers. One key event during the financial crisis was the compound systemic impact of seemingly unrelated individual lender decisions not to lend to borrowers.

Figure 13.1 shows a traditional linear- and event-based view expressed by Bank of England (2007). The global financial market turmoil was triggered by an increased default rate in the US sub-prime mortgage market. In response to this increase in default rates, asset backed securities were downgraded; investors became more risk averse and as a result lost interest in financial instruments that seemed to be exposed to the sub-prime markets. Investor weariness spilled over to the short-term global credit markets and as a result major financial services firms faced increased liquidity risk. Due to the need to provide cover for funds typically available to investors on the short-term credit markets, global financial firms experienced a major deterioration of their balance sheets. The result was a tendency to 'hoard' cash and an increased aversion to risk. This triggered tensions on the inter banking lending market resulting in reduced liquidity and higher inter banking interest rates.

In parallel, complex asset classes experienced continued devaluation leading to an increased need for capital, and this forced some firms to shed assets. For instance, Merrill Lynch (later acquired by Bank of America), a US-based financial services company, sold 30 billion dollars worth of assets in July–August 2008. As a result, asset valuations decreased and firms needed yet more capital. At the same time credit agencies started to revise their risk ratings (i.e., risks were perceived as being much higher) and individual firms therefore required more capital to cover potential losses. At this point the procyclical forces outlined by the FSB report were in full play and this, as we all know now, ultimately had a very negative impact on all major developed economies.

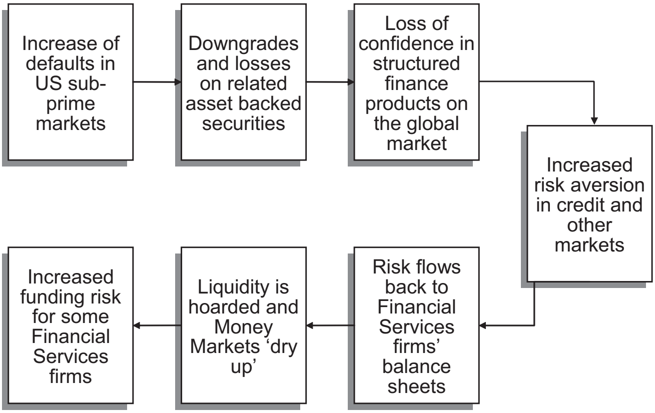
Increase of defaults in US sub-prime markets
Downgrades and losses on related asset backed securities
Loss of confidence in structured finance products on the global market
Increased risk aversion in credit and other markets
Risk flows back to Financial Services firms' balance sheets
Liquidity is hoarded and Money Markets 'dry up'
Increased funding risk for some Financial Services firms

Figure 13.1 Crisis 'Phases' of the 2007 Financial Markets Turmoil (Based on Bank of England's 2007 Financial Stability Report's Chart 1)

In April 2008, the Financial Stability Forum provided a summary of how various risk management processes 'broke down' during the crisis (p. 16).

- Regulators and supervisory bodies failed to identify the risks associated with financial services firms' structured products and other types of off-balance sheet entities. As a result firms ended by having insufficient capital buffers to deal with asset devaluations and decreased investor risk appetite.
- Financial firms misjudged the risk associated with off-balance sheet entities, often due to an over-reliance on risk ratings provided by credit rating agencies. Instead of making their own analysis, firms relied on ratings provided by specialised companies such as Standard & Poor's.
- Commonly used metrics to assess market risk, that is, '… the risk of losses in on-and off-balance sheet positions arising from movements in market prices' (Gallati, 2003: 34), such as VaR, could not be leveraged. A primary reasons for this was that VaR required historical data and the ability to Mark-to-Market, that is, to assign an asset a value based on the current market price of the same or similar type of asset. But this could not be done because there were no markets for the particular asset type! The reason for this was that risk aversive investors had caused markets to quickly 'dry up'. As a result, valuations were model-based and therefore did no reflect a realistic market value.
- Investors (i.e., market participants) misjudged at least one of the following: a) borrowers risk of defaulting, b) the dangers associated with making too many investments in a single type of assets or financial instrument, and c) the risks associated with reduced liquidity, that is, the ability to turn an asset into cash without any impact on the value.

Central banks, such as the Bank of England, the European Central Bank, and the US Federal Reserve Bank reacted by cutting interest rates and by pouring money into the global Financial Services System. Many national governments also provided funds to key financial institutions with the intention to kick-start the key process between providers and users of funds. In addition, unprecedented efforts began to create macro-prudential processes, operating at a system level in the FSS.

One of the key lessons from the recent crisis is that most risk-management efforts by regulatory bodies typically focused on individual firms and not on understanding systemic risk. In Europe, the de Larosière report proposed to establish a European Systemic Risk Council with a charter to '… form judgments and make recommendations on macro-prudential policy, issue risk warnings, compare observations on macro-economic and prudential developments and give direction on these issues' (de Larosière, 2009: 44). In March 2009, the US Treasury department similarly recommended to form a so-called 'Systemic Risk Regulator' focused on capturing systemic risk. A key capability of both of these proposed systemic regulatory bodies is the need to understand how various components of the FSS are interconnected and how their behaviour can possibly impact the broader economy. In addition, systemic regulatory bodies must be able to assess if the 'sum is larger than its parts' from a risk perspective. In other words, how do risks associated with individual components of a FSS 'add up' or combine? One question is whether they behave linearly at all? Most important of all, the proposed regulatory bodies need to have a shared understanding of what constitutes the 'Financial Services System'.

What is the Financial Services System?

General Systems Theory, developed by the Austrian biologist Ludwig von Bertalanffy, defines a system as '… a complex of elements standing in [dynamic] interaction' (von Bertalanffy, 1975: 159). The primary focus of General Systems Theory is to identify how systems reorganise and self-regulate to achieve their objectives. Von Bertalanffy spent most of his life trying to understand so-called 'open systems', that is, systems that constantly exchange material – matter and energy – with their environments, very much as cells do in a biological organism. Financial Services Systems are open systems that must adapt to a changing environment by means of constant 'exchanges' with their environments. Without these exchanges open systems cannot sustain performance and will ultimately enter a catastrophic state, i.e., an irreversible state of failure (cf., Sundström and Hollnagel, 2006).

Conceptually, exchanges between a FSS and its environment, as well as within the system itself, are a mixture of the following two types:

- *Proactive Exchanges* driven by the FSS's goal, such as the perpetual need to identify demands for its services, to meet the demands by providing financial services, and ideally to do it so that its assets and resources are either preserved and/or increase in value.
- *Reactive Exchanges* that are 'forced' upon the system by the sheer dynamics of the environment. For example, a general reduction in households' disposable income is likely to drive up delinquency rates on loans and thus eventually will force a FSS to take actions, including a write-off of delinquent loan related losses.

Most FSSs are regulated by various types of regulatory bodies such as central banks, or organisations like the UK's Financial Services Authority. The primary goal of these entities is to make sure that any particular FSS is not in a state that results in increased systemic risk. The primary material exchanged between a regulatory entity and a FSS is information. This information is used by regulatory bodies to assess the state of the system, to provide regulatory guidance, and to formulate action plans. Other types of entities, such as credit agencies, are also part of the Financial Services System. A credit agency is basically an organisation dedicated to collecting information about the credit worthiness of individuals and organisations. This information is provided to other entities in the system that often leverage this information to make lending, or, investment decisions.

Figure 13.2 illustrates the following key points about Financial Services Systems (or any system):

- System boundaries are defined relative to a particular perspective. An economist looking at the FSS from a global perspective is likely to include everything in the two white rectangles. A 'Micro Regulator' is likely to focus on the entity it is regulating. For example, in Figure 13.2, 'Micro Regulator A' regulates financial services firm B, whereas 'Macro Regulator A' is responsible for monitoring the overall FSS. The key lesson from Figure 13.2 is that system boundaries will be 'drawn' as needed by the entity looking at the system.
- The types of entities ('institutions') included in a view of a FSS will depend on what attributes and behaviour an entity must display to be considered part of the FSS.

Following von Bertalanffy's definition of a system, we obviously need to answer the question of what creates the dynamic interactions in Financial Services Systems. To accomplish this, we will use the definition that a FSS 'can be defined in general

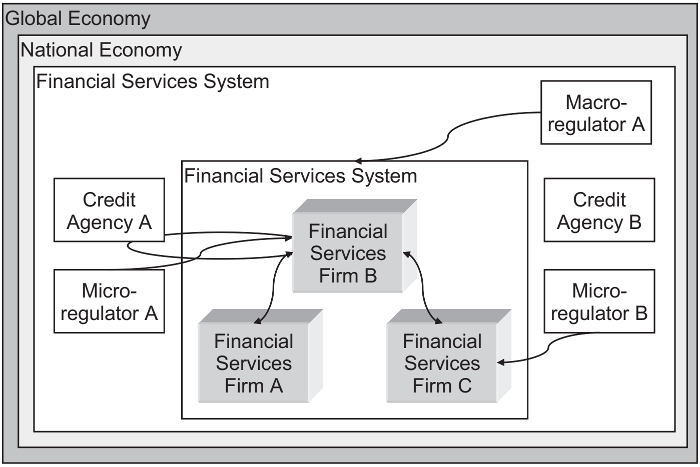
Global Economy
National Economy
Financial Services System
Macro-regulator A
Financial Services System
Credit Agency A
Credit Agency B
Financial Services Firm B
Micro-regulator A
Micro-regulator B
Financial Services Firm A
Financial Services Firm C

Figure 13.2 Different views of financial services systems and entities

terms as the interaction between supply of and the demand of the provision of capital and other finance related resources' (Schmidt and Tyrell, 2004: 21). From this definition we can infer that the interaction of demand and supply functions will have an impact on any FSS.

What Creates the Dynamic Interactions?

The financial services literature often makes a distinction between a so-called institutional and a functional perspective (Merton, 1995). The institutional perspective builds on the structure of existing institutions in the Financial Services industry, whereas a functional perspective '… takes as given the economic functions performed by financial intermediaries and asks what the best institutional structure to perform those functions is' (Merton, 1995). In the present chapter we assume that interactions of economic functions are the primary source of the behaviour of a FSS, both from a micro-and a macro-prudential perspective.

To understand the behaviour associated with the dynamic interactions within and between FSSs, and of course to determine the system boundaries, we need to have a model or a representation of the FSS. Generally speaking, a model is a simplified representation of the salient features of an object system that can be used to analyse or reason about it. Some models can be explicit and visualised while others are implicit, e.g., embedded in a program or a set of equations. A key feature of any model is that it determines what information is relevant and how information is organised and processed. A model will influence how we define system boundaries and make decisions about what we consider to be entities of a particular FSS.

While there is a large variety of models, few are able to address the risks that arise from performance variability, even though this variability may lead to both negative and positive outcomes. The Functional Resonance Analysis Method (FRAM), proposed by Hollnagel (2004) offers a way to understand how functions can be coupled or interconnected. The principle of 'functional resonance' makes use of the definition of stochastic resonance concepts used in physics. Formally defined, stochastic resonance happens when a non-linear input is superimposed on a periodic modulated signal that normally is too weak to be detected, so that a resonance between the weak signal and the stochastic noise pushes the result over the threshold (Hollnagel, 2004: 168).

Stochastic resonance is generally used to illustrate the noise-controlled onset of order in complex systems. The difference between stochastic and functional resonance is that the variability that constitutes the 'noise' in the latter case is systematic, hence predictable, rather than random. From a functional perspective, resonance means that the variability of a set of interconnected functions may combine to affect the normal but otherwise undetectable variability in other functions and thereby lead to significant unexpected outcomes.

While the details of FRAM modelling still are being refined, the basic principles have been established and demonstrated particularly in the aviation domain (e.g., Hollnagel et al., 2008). In the next section, we will highlight the modelling steps, leverage some of them and also introduce the visual notation used by the method.

The Modelling Steps

This section will demonstrate how functions of a FSS can be modelled and how this can help to uncover interdependencies. The primary purpose is to illustrate the modeling process rather than to present a validated model of the functions of a FSS.

A FRAM modelling process typically consists of four phases:

1. Identify what functions need to be modelled.
2. Identify conditions that could lead to change in performance.
3. Identify areas where functional resonance could emerge.
4. Identify how performance variance can be monitored and controlled.

In the following, we will focus on phases 1, 3 and 4.

Identifying the Core Functions

The primary purpose of this modelling step is to identify the critical functions of a system. Merton and Bodie (1995) proposed a functional view of FSSs, cf., also Merton and Bodie (2005), which comprises the following functions:

- *Clearing and Settlement* – payment services related to the exchange of goods and services.
- Risk Management – manage uncertainty and development and implementation of mitigation strategies and action plans.
- *Transfer Economic Resources* – facilitate the flow of economic resources targeting efficient use of capital.
- *Information Sharing* – provide stakeholders with access to risk ratings, price information and other types of information perceived as critical for making decisions in, or about, FSSs.

Following Schmidt and Tyrell (2004), we need to add *Demand and Supply* functions to the four functions mentioned above. Finally, we view the *Resource Pooling* function as being part of a mechanism that can be used to transfer economic resources and as a result combine functions the *Transfer Economic Resources* and *Resource Pooling functions* into one.

After the key functions have been identified, the next step is to generate a high-level description of each function. In the FRAM modelling framework, a function can be described by the following attributes:

- *Input (I)*: that which the function processes or transforms, or that which starts the function. In financial services, this could be the information that is modified, interpreted, or used in any other way by the function.

- *Output (O)*: that which is the result of the function, either an entity or a state change. In financial services, this can be economic resources for a party that previously did not have any economic resources. For example, Mr. Smith goes to a bank to apply for a mortgage and the bank decides to transfer economic resources to Mr. Smith, that is, provide Mr. Smith with a loan of some type.
- *Preconditions (P)*: conditions that must be exist before a function can be executed. An example in financial services is the existence of a financial market with a defined demand and supply.
- *Resources (R)*: that which the function needs, or consumes, to produce the output. This could be some type of financial assets and/or market participants such as investors.
- *Time (T)*: temporal constraints affecting the function (with regard to starting time, finishing time, or duration). For example, the time of a transaction can greatly influence the value of an output in the financial services industry.
- *Control (C)*: how the function is monitored or controlled. In financial services this can be a firm's risk management function and /or a regulatory function such as a Central Bank. Table 13.1 lists the instantiations of the attributes of the two functions Risk Management and Transfer Economic Resources for Firm A.

In Table 13.1, the values of the attributes are shown as simple labels. Each value can, however, be described as the output of a function and the description of the initial functions can therefore be expanded until all the necessary functions have been defined. Notice also that the actual 'mechanisms' associated with execution of the functions do not need to be specified at this stage.

Table 13.1 also shows how two functions can be coupled by means of shared attribute values. For example, *Risk Assessment* is an output of 'Risk Management' as well as a pre-condition for 'Transfer Economic Resources'. Functions can also be coupled by means of common attribute values, such as *Regulators* and *Management* that are part of the control attribute for both functions.

Table 13.1 FRAM descriptions of the functions 'Risk Management' and 'Transfer Economic Resources'

Firm A: Risk management	
Attribute	**Assignments/values**
Input (I)	Risk related data; Risk ratings; Risk 'appetite'
Output (O)	Risk assessment; Decisions
Preconditions (P)	Required data available; Perceived need to manage risks
Resources (R)	Risk methodology/processes
Time (T)	Deadline for risk assessment
Control (C)	*Regulators; Management*; Risk policies and procedures
Firm A: Transfer economic resources	
Attribute	**Assignments / values**
Input (I)	Request for economic resources
Output (O)	Economic resources available to requestor
Preconditions (P)	Risk assessment; Transfer mechanisms
Resources (R)	Funds
Time (T)	Deadline for transfer
Control (C)	Regulators; Management; Transfer policies and procedures

Identifying Potential for Functional Resonance

Functional resonance is an emergent attribute of a system, which means that it cannot simply be derived from descriptions of the constituent parts of the system. Unlike the traditional safety paradigm, Resilience Engineering does not assume that there is a simple causal relation between the parts and the whole. In a FSS, a single entity with a high default risk (i.e., the risk that a borrower will not be able to meet debt obligations) may be coupled to other entities with a high default risk via a securitisation process. When otherwise independent entities become interconnected this can create the potential of functional resonance in the overall system, that is, an amplification of the impact of outcomes that may lead to disproportionate losses (or gains). The output of one function may, for instance, provide the resources of another function. Or, the output of one function may be a pre-condition for another function. Table 13.2 lists the values that are used to describe the attributes of the two functions 'Demand Economic Resources' and

'Supply Economic Resources'. Figure 13.3 provides a graphical illustration of the most important couplings between the functions described in Tables 13.1 and 13.2.

Some key lessons from Figure 13.3 include:

- If Firm A completely depends on markets to generate its 'raw material' (funds) in order to be able to transfer economic resources, any unusual variation in the system level functions will have impact on Firm A's ability to meet customers' requests.
- While Firm A might be able to generate funds to fulfil customer requests, the ability to sustain performance depends on how Firm A can replenish the 'raw material', that is, capital. One way for Firm A to replenish funds is by having a traditional banking function. Another method is to leverage capital markets.
- The *Demand and Supply* functions are interconnected, that is, *Supply* requires a certain *Demand* and vice versa. The interaction between these two functions is the key subject for economists interested in the laws of demand and supply. From a functional modelling perspective, the focus should be on how the behaviour of the two functions impact individual firms. From a systemic perspective, the impact of financial services' demand and supply functions on the overall system should be the primary interest of economists and macro-prudential regulatory bodies such as central banks.

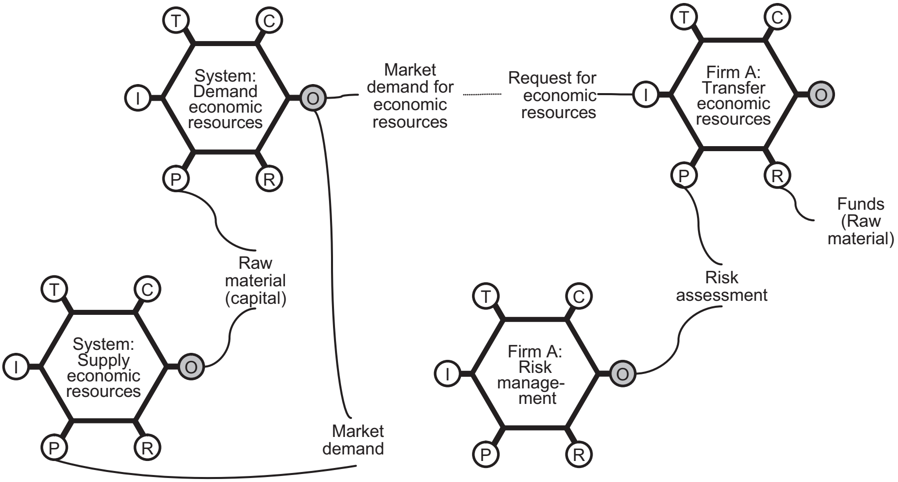
T
C
I
O
P
R
System: Demand economic resources
Market demand for economic resources
Request for economic resources
Firm A: Transfer economic resources
Funds (Raw material)
Raw material (capital)
System: Supply economic resources
Risk assessment
Firm A: Risk manage- ment
Market demand

Figure 13.3 Interdependency of Firm A's risk management and transfer of economic resources functions

Table 13.2 Two FRAM functions with assigned values

System: Demand Economic Resources	
Attribute	**Assignments/values**
Input (I)	Cost of credit; Preference for economic resources type
Output (O)	Market demand for economic resource
Preconditions (P)	Raw material (capital); Potential 'buyers'
Resources (R)	Not defined
Time (T)	Deadline for risk assessment
Control (C)	Regulators; Management
System: Supply Economic Resources	
Attribute	**Assignments/values**
Input (I)	Cost to produce economic resource
Output (O)	Raw material (capital)
Preconditions (P)	Knowledge and tools to establish resources; Market demand for economic resource
Resources (R)	'Raw' material (capital)
Time (T)	Request fulfilment deadline
Control (C)	Regulators; Management

Identifying How Performance Variance can be Monitored and Controlled

There are three types of management and control functions in a FSS. First, an individual firm's management function responsible for driving business results while managing risks at the same time. Second, a micro-prudential regulatory function responsible for oversight of individual firms. Third, a macro-prudential regulatory function responsible for oversight of the overall FSS. In order to be able to monitor and, as required, control and regulate performance of individual firms, groups of firms and of course the overall FSS, each of these functions needs the following:

- A view of the system that is being monitored.
- A view of what needs to be monitored.
- Data and metrics required to perform the monitoring.

A functional model can provide guidance with respect to the first two. A risk assessment and monitoring methodology may provide guidance for the third. We can use a functional model to consider how a micro-prudential regulator might choose to monitor Firm A, based on the functional interdependencies illustrated in Figure 13.3. Table 13.3 lists the values that are used to describe the attributes of how a micro-prudential regulator may choose to monitor Firm A.

The description of the main function of a micro-prudential regulator in Table 13.3 shows the following:

- A micro-prudential regulator cannot monitor and/or determine Firm A's risk profile without including data about system level components such as the Demand and Supply functions.
- A micro-prudential regulator needs to know what data is required to perform risk assessments, that is, the regulator needs to understand Firm A's business model and services and have a risk assessment and monitoring methodology that helps to specify what data is required.

- The Information Sharing function is a critical component of the overall Financial Services System. Without transparency, various stakeholders are likely to lose trust.

Table 13.3 The FRAM functions of a micro-prudential regulator (of Firm A)

Micro-prudential Regulator: Firm A	
Attribute	**Assignments/values**
Input (I)	Firm A's credit portfolio risk; Risk rating; Firm A's risk 'appetite'; Demand/supply of economic resources; Firm A's business model; Risk profile of comparable firms
Output (O)	Firm A's risk profile; Decisions; Guidance
Preconditions (P)	Perceived need to assess Firm A's risk profile
Resources (R)	Risk methodology & processes; Data
Time (T)	Per examination schedule; *ad hoc* based on (P)
Control (C)	Regulators; Risk policies and procedures
Micro-prudential Regulator: Information Sharing	
Attribute	**Assignments/values**
Input (I)	Information sharing request
Output (O)	Economic resources available to requestor
Preconditions (P)	Required data available (from overall system and Firm A)
Resources (R)	Data from Firm A; Processing mechanism
Time (T)	Per examination schedule; *ad hoc*
Control (C)	Regulators; Firm A's management

The next section will use the above principles to show why Northern Rock's leaders and regulators failed to have a valid view of Northern Rock and its interdependencies with other system level components.

Example: The Demise of Northern Rock

Northern Rock was listed on the London Stock Exchange in 1997 and was seen as a very successful financial services firm.

All of this changed when on 14 September, 2007 Northern Rock had to ask the Bank of England for a line of credit to overcome a 'liquidity crunch'. The damage to Northern Rock's reputation was lethal and the firm experienced a classic bank run. In February 2008, Northern Rock was nationalised and as a result is no longer traded as a public company. In retrospect, Northern Rock was one of the first institutions outside the US that experienced the impact of the 2007–2009 crisis in the global financial system.

Northern Rock's business model is illustrated in Figure 13.4. Basically, Northern Rock's business consisted of transfer of economic resources to retail customers. Northern Rock's ability to transfer economic resources depended on two sources of funding:

- Twenty-five per cent of customers' requests for mortgages were met by bank deposits, that is, funds from banking customers' deposits.
- The remaining 75 per cent were fulfilled by getting funds from the global credit markets. As we all know these markets started to experience major distress in the summer of 2007.

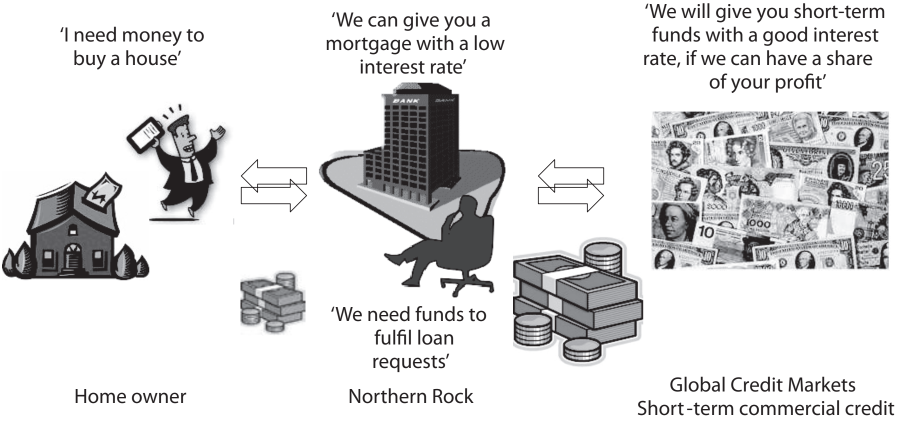
'I need money to buy a house'
'We can give you a mortgage with a low interest rate'
'We will give you short-term funds with a good interest rate, if we can have a share of your profit'
BANK
'We need funds to fulfil loan requests'
Home owner
Northern Rock
Global Credit Markets
Short-term commercial credit

Figure 13.4 Northern Rock's business model

The simplified representation in Figure 13.5 illustrates the risks that were a consequence of Northern Rock's reliance on funding from global credit markets to fulfil 75 per cent of the mortgages. (Grey hexagons indicate that a function operates outside of Northern Rock's control.) This funding depended on a consistently positive output from the Global Credit Markets' 'Transfer Economic Resources' function. However, key preconditions for this were good credit ratings and positive Investor risk assessments of the financial instruments offered by Northern Rock.

When financial markets became increasingly distressed in 2007, investors' trust in mortgage backed financial instruments severely deteriorated resulting in zero demand. As a result Northern Rock experienced an acute liquidity crisis due to its high dependence on mortgaged backed securitisation as a method of creating funds to fulfil loan requests by their customers.

Concluding Remarks

The primary purpose of the present chapter was to illustrate how Resilience Engineering can provide the financial services industry with a different way to understand risk at both a macro-

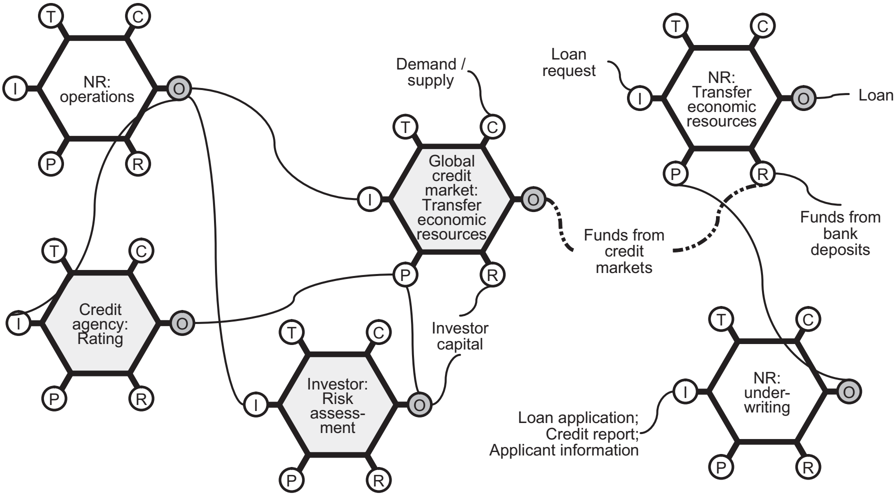

NR: operations
T
C
I
O
P
R
Demand / supply
Loan request
Global credit market: Transfer economic resources
NR: Transfer economic resources
Loan
Funds from credit markets
Funds from bank deposits
Credit agency: Rating
Investor capital
Investor: Risk assess-ment
NR: under-writing
Loan application; Credit report; Applicant information

Figure 13.5 Functional view of Northern Rock's risk exposure

and micro-prudential level. We discussed the need of the various stakeholders to be clear about what the FSS is in any particular situation. In addition, multiple stakeholders will need a shared view of the system. As discussed in the chapter, creating such a shared view requires a common model of the FSS. In this chapter, we suggested to use a functional approach leveraging the work by Merton (1995) and others. Such an approach enables stakeholders and decision makers to focus on behaviour rather than on the specifics of an individual financial services institution. A key advantage of a functional perspective is that it becomes possible to discover risks created by functional interdependencies among individual institutions and system components, as illustrated by the case of Northern Rock. The concepts and functional modelling approach outlined in the present chapter thus provides the basis for the development of a standardised method to capture and better understand risk in the financial services industry.

PART IV
Dealing with the Factual

This page has been left blank intentionally

Chapter 14
To Learn or Not to Learn, that is the Question

Erik Hollnagel

The last of the four main capabilities of resilience – the fourth cornerstone – is the ability to learn from the past. It only comes last, however, because the four have to be listed one after the other. In practice, the ability to learn is just as important as the ability to respond, monitor and anticipate, as discussed in the Epilogue. A system cannot be called resilient if any one of them is missing and the absence of one cannot be compensated for by the increased quality or quantity of any of the others.

The Conditions for Learning

When the importance of learning for safety is considered, it is generally taken for granted that the basis for learning must be things that have gone wrong, such as incidents, accidents and catastrophes. This obviously makes sense from a classical safety perspective since knowledge of why things have gone wrong in the past is essential to prepare for the future. To learn from things that have gone wrong is, however, less reasonable when considered from a learning perspective, since it clashes with the basic principles for effective learning.

In order for learning to take place, three conditions must be fulfilled. The first condition is that there are reasonable opportunities to learn, that is, that situations where something can be learned occur with a sufficiently high frequency. (More precisely, the situations must occur so often that the lessons learned from previous situations have not been half-forgotten.) The second condition is that the situations are sufficiently similar to allow generalisations to be made, that is, they must have something in common or be comparable in some sense. People and organisations must be able to recognise that there is something in situation A that also can be found in situation B, not just in the manifestation of the outcomes but also in the reasons or causes. Finally, the third condition is that there must be sufficient opportunity to verify that the right lessons have been learned. This can be seen as a kind of combination of the first and second condition, in the sense that a comparable event must happen before the lessons have been forgotten – and hopefully well before the lessons have to be used in an actual event.

If we consider accidents, or even more extreme events such as emergencies and catastrophes, it is clearly important to find out why they happened but also clear that they do not offer the best basis for learning. Accidents do not happen very frequently, at least if the domain of activity is reasonably safe. Accidents are furthermore usually different from each other, and the differences are often proportional to the magnitude of the outcomes. Lastly, because accidents happen so rarely, there is little opportunity to check whether the right lessons have been learned. Accidents therefore do not provide good conditions to learning, common stereotypes notwithstanding.

It follows from these arguments that learning can be more effective if it is based on events or conditions that happen more frequently and that – almost by virtue of that fact – are less extreme and less dissimilar (e.g., Herrera et al., 2009; Woods and Sarter, 2000). Indeed, it is more efficient to learn from what goes right than to learn from what goes wrong (cf., the Prologue), because the former happens far more often than the latter. This position is also consistent with the basic principle of Resilience Engineering that failures are the flip side of successes and that they both have their origin in performance variability on the individual and systemic levels (Hollnagel et al., 2006: xi).

The Impact of Learning

The four main capabilities of resilience are equally necessary and therefore equally important. Taking learning as a starting point, it is easy to argue that the ability to respond would be of little value without the ability to learn. A nation, a system, an organisation or an individual can, of course, adopt a set of predefined or stereotypical responses – and sometimes do. As long as the characteristics of the environment do not change, as long as nothing unexpected happens, the set of responses may be adequate. But unless the environment remains stable, or unless the environment can be completely controlled so that nothing unexpected happens, the pre-defined responses will sooner or later become inadequate. In other words, it will be necessary to learn new ways to respond, as discussed in Chapter 1. And it is by observing and evaluating the efficiency of the responses that the system (nation, organisation, individual) can learn.

A similar kind of argument can be made for the relation between learning and monitoring. The choice of which indicators to monitor can be based on formal or *a priori* models but it is 'rare to find explicit models that provide a formal basis for identifying measures', as argued in Chapter 5. It is mainly by learning through practice that the proper basis for monitoring – the indicators that must be watched – can be established. Yet simply adding new indicators whenever something has happened is not efficient in the long run. (A handy illustration of that is the way that most anti-virus software relies on a list of virus signature definition.) The efficiency of monitoring depends on the efficiency of learning, just as the efficiency of learning depends on looking at the right kinds of experiences.

Finally, learning is also necessary for anticipation. As Chapter 9 points out, one of the patterns of anticipation is that 'resilient systems are able to recognise the need to learn new ways to adapt'. In relation to anticipation, learning is essential to produce a realistic, or even adequate, model or understanding of what may possibly happen in the future. This highlights the importance not only of learning, but of learning the right lessons, that is, of understanding what has happened in a way that is useful for the future functioning of the system. The worlds of industry, to say nothing of the world of politics, business, and finance, provide ample examples of how difficult this is (Touchman, 1985).

What Should Be Learned?

Learning is not a mechanical function, and cannot be reduced to data collection or statistics. For every use of learning it is crucial to learn 'the right thing.' But what exactly does that mean?

In relation to safety it is frequently pointed out that it is important to be thorough in learning. Given that learning traditionally has been based on things that have gone wrong, one advice has been to look for second stories beneath the surface (Woods and Cook, 2002). Unlike the traditional search for 'root causes', this does not mean that one should go as far back as possible, but rather than one should consider possible alternative explanations. Another advice has been captured by the phrase 'what-you-look-for-is-what-you-find' (WYLFIWYF) (Lundberg et al., 2009). The WYLFIWYF principle means that explanations of accidents are strongly influenced by assumptions about how different factors interact, that is, the accident ''mechanisms.' And since it is impossible to learn what has not been found, the corollary of the principle is that 'what-you-find-is-what-you-learn' (WYFIWYL).

The four chapters in this section each contribute valuable advice on how learning can be improved. Chapter 15 focuses on the obvious importance that the factual functioning of the system (where the events took place) must be transparent. This means that it is important to gather evidence about how the system functions, over and above looking for direct causes. In cases where things have gone wrong (accidents) or may go wrong (risk assessment), this is relatively straightforward and is often practised as an institutionalised process. In other cases, for instance when it is uncertain whether an organisation is sufficiently safe, the gathering of evidence is more difficult and may require a protracted process and investment of time and effort. Yet without the evidence, the situation cannot be properly assessed, and without that learning cannot take place.

Chapter 16 emphasises the importance of going beneath the 'surface', by illustrating how things may go wrong because of communication failures. It is argued that it is crucial to study coordination mechanisms in order to understand the resilience of socio-technical systems. This is therefore also an argument for extending the study beyond accidents and failures. The ways in which various entities of an organisation, in this case a health-care system, coordinate their activities and in particular how practitioners adapt to the unexpected using emergent coordination mechanisms, can most easily be seen in cases where the coordination succeeds and where therefore nothing goes wrong. Coordination is a function rather than a failure, but it is a function that can and must vary to match the working conditions.

Chapter 17 takes a closer look at the way data can be obtained, in this case from incidents. Incident reporting is not something that can be established by an edict, and successful incident reporting requires more than a simple set of tools or procedures. One issue is the pass criterion, that is, how easy it is to distinguish between important and negligible events. A second is the degree of standardisation or regularity of the work situation, since this determines how much information it is necessary to gather. A third issue is the visibility of events, and the question of whether sharp-end operators always are the best source of information. A fourth is how the heterogeneity of the community affects the scale of the reporting system. And a fifth issue is the current safety culture, which may determine whether reporting can be anonymous or not. All five issues are important for the basis of learning, whether it is narrow or extensive, how subjective and reliable it is, how complete it is, and when incident reporting will work or when it will not.

Chapter 18, finally, takes a closer look at how differences in cultures and occupations may affect what is considered safe and what is considered risky, hence affect what is learned. The chapter shows that there are clear variations in perspectives according to the national culture and occupation of respondents. People from different cultures and different occupations may therefore not learn the same lessons from what ostensibly are the 'same' events. This variety in outcomes is, however, not necessarily a weakness but may, following the arguments of Woods and Cook (2002) in fact be a strength, if only the organisation is able to make use of it.

Chapter 15
No Facts, No Glory

John Stoop

In order to learn from the past in complex, dynamic systems their factual functioning must be transparent. Knowledge deficiencies and systemic deficiencies should be derived from analysing decision-making and underlying assumptions, uncertainty and knowledge. Such deficiencies are to be identified throughout the design process as well as during operational practices, in order to facilitate systems change and enhancement of the safety performance. Safety investigations should facilitate evidence-based learning. Without a fact-finding mission dealing with on-site observations and interpretations, evidence-based learning will be lacking and achieving consensus will be reduced to a mere negotiation result. Providing evidence is the domain of forensic sciences. By providing a timely transparency, safety investigations may take the role of providing functional requirements for the (re-)engineering design of resilience as a system property. This approach is demonstrated by two major case studies that have been conducted in the Netherlands: the El-Al air crash in Amsterdam and the inquiry into the High Speed Line ERTMS signalling system.

Introduction

According to the methodological principles for engineering design, systems are assessed with regard to their intended functioning throughout the design process (Stoop, 1990). Historically, technological aspects of safety have been identified at the three consecutive phases of conceptualisation, function allocation and materialisation focusing on, respectively:

- *Inherent design principles*, to be deployed in selecting specific concepts, configurations and control arrangements dealing with fail-safe, crash worthiness and delegated or distributed responsibilities in human centred control.
- *Emergent system properties*, described by the acceptability of accident scenarios in their operational context in fulfilling their designed functions. Such properties are identified as redundancy, robustness, reliability, reconfiguration, rescue, recovery and resilience.
- *Performance indicators*, by quantification of the *probability of operational failures* and their consequences as either expressed in safety standards, accident frequency rates or safety integrity levels.

Determination of systemic and knowledge deficiencies originating from these phases do not, however, restrict themselves to the technical investigations of single events in the operational phase, commissioned to accident investigation committees in case of a major occurrence. As the HSL/ERTMS and Schiphol/B747 case studies indicate, investigations also can be applied to the level of complex operational management in a multi-actor environment and the design and development phase of a major project (Stoop et al., 2007; Stoop, 2009). The initiating organisation in both investigations was the Dutch Parliament, requesting transparency in the event and underlying decision-making processes.

Accident investigation cannot evade the challenge of dealing with the notion of 'cause'. Accident investigations have to provide evidence, based on scientific methods, principles and knowledge (Stoop, 1997; Sklet, 2002). By their mandate, accident investigations must render advisory opinions to assist the resolution of disputes affecting life or property. They play a role as public safety assessor in a multi-actor environment in open decision-making processes. During their fact-finding phases, such investigations primarily rely on forensic sciences.

In general, forensic sciences comprise the science, methodology, professional practices and principles involved in diagnosing common types of accidents and failures. Determination of causes of technical failure and gaining oversight over social system performance requires:

- familiarity with a broad range of disciplines
- The ability to pursue several lines of investigation simultaneously.

As such, these are essential skills of accident investigators, engineers, researchers and managers who are dealing with safety investigations.

Safety investigations can be seen in analogy to forensic fact-finding missions as providing learning potential for understanding complex and dynamic systems. Investigations provide a transparency of design and operational decisions that cannot be replaced by management oversight, audits or governmental inspections during the operational phase of a system. Investigations may identify the nature of design and operational uncertainties.

Case 1: Before, During and After the Event; the Boeing 747 Case Study

On October 4 1992, a B-747 freighter of El-Al crashed into an apartment block in Amsterdam, killing four people on board and 43 people on the ground. During its outbound flight, two engines separated from the airplane, rendering the aircraft uncontrollable after which it crashed into the Bijlmermeer, a suburb of Amsterdam. During the investigation, the direct cause of the accident was established as a design flaw in the engine pylon, due to which the mounting of the fuse pins failed. However, the consequences of the crash far exceeded the direct causes and the technical lessons learned on the short term.

In the survey three main phases are discriminated.

- Elaborating the design decisions and the role of incremental change in extrapolating the Boeing 707 engine mounting concept into the Boeing 747 design.

- The actual investigation with respect to its complexity and case specific implication on a national policy making level over a period of about seven years, including a parliamentary inquiry.
- The implications of the crash for the international aviation community on the long term.

The benchmark nature of this disaster is in the wide implications it has had on the safety assessment of full freighters, the establishment of a multi-modal independent accident investigation agency in the Netherlands, public risk perception with respect to external safety and the shift towards a focus on dealing with the aftermath of a disaster and the role of public governance in crisis policy decision-making.

The Reason for Building such Aircraft

The Boeing 707 was developed in the early 1950s and fitted with four turbojets. In the 1960s it was the common passenger aircraft. It established Boeing as one of the largest manufacturers of passenger aircraft and led to later series of aircraft with 7*x*7 designations. The huge growth in air travel made the 707 a victim of its own success. The 707 was too small to handle the increased passenger densities on the routes for which it was designed. The solution was to design an aircraft with 2.5 times the size of the Boeing 707. The production of the 747 was under enormous pressure to succeed. This was because there was an Anti-Jumbo Lobby, which questioned the safety of the big aircraft (Hoogendorp, 2007).

The design of the engine nacelle and pylon incorporated provisions that prevent a wing fuel cell rupture in case of engine separation. This prevention was achieved with structural fuses.

During the design of the strut-to-wing attachments of the 747, Boeing employed a fuse pin concept that was similar to the 707. This was designed to be fail-safe for vertical loads. The concept of 'safe separation' was addressed by mounting a fuse pin in each of the attachments, preventing this catastrophic damage.

Several related accidents with the 707 occurred. Before the Bijlmermeer aviation disaster with the 747, there were four related accidents with the 707. The separation of the engines in the incidents was attributed completely to external forces acting on the engine. Consequently, the industry had no experience with unsafe in-flight separation of the engine due to mechanical defects. In April 1988, Boeing received a report of a crack in a new style fuse pin. In total, 14 instances of cracks in new style mid-spar fuse pins and 9 reports of cracks in new style diagonal brace-fuse pins were reported. With the 747 several related accidents also occurred. In total there have been five accidents with separation of an engine assembly on the 747. As a consequence, the design and certification of the Boeing 747 strut was determined to be inadequate to provide the required level of safety. The new design philosophy:

- designs for strength, durability and maintainability
- employs the damage tolerance design philosophy
- addresses the possibility of damage due to heat, corrosion and accidents
- addresses ultimate load conditions that are not all inclusive, such as multiple blade-out conditions, multiple ultimate load conditions and unusual turbulence.

It can be concluded that Boeing had knowledge about the weakness of the construction before the Bijlmermeer disaster. After some accidents, with and without fatalities, Boeing developed a new design for the engine mounting.

The Bijlmermeer Crash

After the crash of the 747 in the Bijlmermeer and the other related accidents with the 747, Boeing designed a new system for the engine attachment. This has been applied to all new types. At this moment, the Boeing 747-400 is the only model still in production.

Immediately after the crash, several activities were undertaken by the Dutch government.

The NASB Investigation into the Crash

In its conclusions the Netherlands Aviation Safety Board (NASB) established the Probable Causes referring to the design deficiencies of the B747 (NASB, 1994):

> The design and certification of the B 747 pylon was found to be inadequate to provide the required level of safety. Furthermore the system to ensure structural integrity by inspection failed. This ultimately caused – probably initiated by fatigue in the inboard midspar fuse-pin – the no. 3 pylon and engine to separate from the wing in such a way that the no. 4 pylon and engine were torn off, part of the leading edge of the wing was damaged and the use of several systems was lost or limited. This subsequently left the flight crew with very limited control of the airplane. Because of the marginal controllability a safe landing became highly improbable, if not virtually impossible.

In its 14 recommendations, the NASB referred to the necessity to redesign the pylon structure, including full-scale fatigue and fail-safe testing, a need for incorporating a fail-safe analysis in the certification procedures, and a need for improved airworthiness measures and associated inspection systems. The recommendations also addressed a joint emergency handling between flight crew and ATC in identifying the severity of the event. An improvement in recovery potential was recommended taking into account not only the safety of airplane and passengers, but also the risk to third parties especially where residential areas should be considered.

The RAND Report

Within weeks after the crash, the Dutch Ministry of Transport, Public Works and Water Management commissioned a safety evaluation of Schiphol Airport about the risk to third parties on the ground in case of an accident (RAND, 1993) The report was issued shortly afterwards, focusing on the modelling and calculation of third party risk as a part of the national safety debate on low-probability/high-consequence events. The evaluation used an analytical approach to the subject due to the lack of empirical data on third party risk, the vast amount of uncertainties involved in crash rate data, the consequence assumptions and the inability to predict the timeliness and effectiveness of enhancement measures.

The report questioned the imposition of single standards for aviation risk exposure which were advocated in analogy with the Dutch regulations on maximal acceptable levels of exposure to toxic substances and hazardous activities.

The evaluation put the development of Schiphol in an international aviation perspective.

- Safety considerations may change as Schiphol evolves into a mainport. The projected growth (2.7 times passengers and 4.5 times freight tonnage over the 1993 situation) would increase third-party risk if a linear extrapolation model was applied. However, mitigating factors such as modern airplanes, technological improvement in ATC and aircraft avionics, additional runway capacity and improved international control over risky airlines may have a positive influence on the actual risk performance levels.
- Safety is an airport-wide problem. Coordination of safety was dealt with informally across various operating organisations within the airport and national government. Since no 'magic bullets' exist, an integrated safety management system was required in order to coordinate, monitor and assess safety procedures of the various actors and stakeholders.

Public Concern about Airport Safety

Public dissatisfaction with the findings and governmental responses eventually lead to a Parliamentary Inquiry in February 1999, shortly before parliamentary elections (Parliamentary Inquiry, 1999). This inquiry was unique in two ways. It was the first time such an inquiry had to deal with a wide variety of stakeholders and actors on a national as well as an international level. At the international level, the relation with Israel was put under pressure due to the preferential treatment of El-Al at Schiphol airport. The inquiry also directly involved the public, and the commission had to deal with the emotions of victims, relatives, rescue and emergency workers and residents of the Bijlmermeer. The population of this suburb represented a wide variety of nationalities, religions and races, each with their own culture, social status and backgrounds.

While the inquiry was of little political significance, on another level it put an end to allegations about poor governance on the cargo manifest and all existing conspiracy theories were rejected by a lack of proof. The public anxiety on the crash was released (Houtman, 2008).

Integral Safety, a Substantive Assessment

In the RAND report, the crash and the expansion of Schiphol airport were considered two critical aspects in aviation risk communication and public decision-making on the expansion of Schiphol airport into a mainport in the international aviation industry (RAND, 1993).

On one hand, regional economic development and urban planning deal with concepts such as compact cities and multifunctional use of space. On the other, transportation corridors in Trans European networks, High Speed Lines and Mainport developments and interoperability demands are addressed from a European Commission perspective. In compliance with the policy of Quantitative Risk Analysis that dominated at the time, and in order to cope with the Post-Seveso Directive of the European Commission, the Dutch Ministry of Internal Affairs has taken up its responsibilities by developing the concept of 'Critical Size Events', defining the nature, size and duration of major incidents in order to accommodate the required resources for rescue and emergency handling (Stoop, 2003). The absence of a mandatory application of such a resource assessment has, however, led to questions on rescue and emergency performance levels and efficiency every time a major event occurs as the recent crash of Turkish Airlines at Schiphol in February 2009 has demonstrated.

In the RAND report, recommendations were made to enforce high safety performance standards at the airport level to maintain the public confidence while facilitating the growth in air travel. To protect the interest of European citizens living in the vicinity of airports or travelling on board third country aircraft, the need to enforce safety standards became apparent (Roelen, 2008). At the international level EASA took the initiative to inspect third country aircraft on the spot. This inspection however, cannot substitute the responsibilities of a country to perform proper regulatory oversight. The Safety of Foreign Aircraft Programme was initiated to perform ramp inspections. Non-complying aircraft and airlines could be banned from entering the EU airspace by putting them on a Black List. Such blacklisting has occurred frequently and updates of the Black List are published on a regular basis.

At the national level, the RAND report recommended to install an Integral Safety Management System (ISMS) at the level of the Schiphol airport community. After its foundation in 1993 this ISMS operated in full transparency and openness with an independent status. The Minister of Transport, however, decided to abolish the safety advisory commission in 2006, replacing this expert commission by the Alders Roundtable (PRC, 2006; Ministerie van Verkeer en Waterstaat, 2008). In this roundtable, all stakeholders – residents, local and regional governance, ministries and the aviation sector – came together to achieve a common advice on the development of Schiphol and the region until 2020. The roundtable addressed selective growth combined with noise abatement, while safety was not included in its mandate.

Lessons Learned

Some lessons can be learned from this survey with respect to using accident investigation as an input for Resilience Engineering. This survey on NSAB and RAND investigations revealed some strengths:

- The substantive independence of the investigation. While the NSAB already had its independence assured, the Schiphol Group had to finance the RAND investigations, but had no say in the report, which was sent to the Ministry of Transport directly.

- The prominent role for single event investigations in learning processes to provide factual information on the functioning of the aviation system in practice. Retrospective learning loops may provide valuable factual information and transparency at the level of the airport community (Stoop and Dekker, 2007). Independent investigations identify learning at the institutional level in order to sustainably incorporate safety at a national as well as international safety policy decision-making level.
- Safety at Schiphol has been broadened from aircraft safety towards an integral safety notion incorporating aircraft safety, airport safety management, external safety and rescue and emergency services.
- The focus of the NSAB successor, the Dutch Transportation Safety Board, has enlarged its scope from investigating before and during, towards after the event, and multi-actor involvement at all systems levels, including governance and control at the higher systems levels.
- Institutional independence of a public safety assessor at the airport level is indispensable in order to facilitate and maintain a sustainable, proactive and integral assessment of operations. The functioning of such an assessor should be assured by legal and organisational arrangements.

However, the survey also revealed some weaknesses:

- It is difficult to maintain transparency at the operational level of an airport community. Safety is balanced against other dominant aspects, such as noise abatement and limitations to growth, while safety awareness fades away some time after an accident.
- Safety policy making is fragmented and revolves across policy domains in time, from crash investigation (Ministry of Transport), via rescue and emergency (Ministry of Internal Affairs) to external safety (Ministry of Housing, Spatial Planning and Environment).

- In such a process, a substantive safety assessment approach is replaced by a consensus achieving decision-making approach. Throughout such revolving, safety gradually reduces from a strategic decision-making criterion to a stakeholder based operational constraint.
- Despite the analytical robustness of both investigations major uncertainties remain for a proactive approach to safety at the airport community level, due to the very low probability and limited availability of event data of a historical nature. Rescue, recovery and resilience abilities after such an unpreventable event remain indispensable.

Case 2: ERTMS. An Inquiry into the Safety Architecture of High Speed Train Safety

The European Rail Traffic Management System (ERTMS) is a part of the renovation and upgrading of national railway systems, facilitating interoperability on the EU rail network and fully software controlled train surveillance. ERTMS is a trend shift from technical compatibility across nations towards standardisation and harmonisation on the main EU network corridors. The Dutch High Speed Line is part of Paris-Köln-Amsterdam-London corridor.

For the Dutch ERTMS development several political choices have been made.

- Innovation in Public-Private Partnerships in contracting, mixing public and private interests.
- The development and implementation of technology is done concurrently instead of sequentially. A simultaneous technical development of standards and software components assumes a pragmatic off-the-shelf availability of components from various industrial consortia.
- Based on cost-effectiveness considerations, no redundancy was incorporated to prevent failure during the transition towards these new technologies, although they occurred with respect to the signaling level migration, software version upgrading and full scale testing of the system.

With the increase in traffic intensity, the system is loaded to its design limits. The fault tolerance in hierarchical systems decreases quadratically with intensity. About its saturation point, the traffic flow becomes unstable. Due to the dynamic feedback of multiple operators – train drivers as well as traffic controllers – operator induced oscillation becomes possible; their fault handling may cause abrupt and progressive collapse of the overall system. To avoid initiating disturbances, an even stricter task performance of the train driver is required. Increasing the punctuality of the time table under high traffic intensity conditions demands an increasing control effort by the traffic controller and train drivers. This aggravates the tactical and operational cognitive workload of the traffic control centres that are forced to communicate simultaneously with several train drivers. Eventually a gridlock situation occurs where all traffic operations must be terminated by a fail-safe system breakdown followed by a gradual and safety critical reset. The underlying organisational mechanisms which threaten public values such as safety versus private business values, can be identified as coping behaviour in order to deal with conflicting values (Steenhuisen and Van Eeten, 2008). Will it suffice to discipline organisations with advanced contracts and incentives, piling up requirements without clarifying inevitable trade-offs?

Two principal strategies are applied to overcome design limitations.

- A process approach. Recognition of value conflicts and subsequently a structuring of the process of communication, coordination, and cooperation among all stakeholders in their decision-making processes, coping between quantifiable private performance indicators and qualitative public values.
- A technological approach. Elimination of the human involvement in disguised bad performance due to ambiguous and hybrid decision-making values by developing an innovative train control system, based on modern technology and a new generation of signalling systems.

Emergent Properties

During the High Speed Line investigations, several value judgements became visible dealing with the project organisation and technological scope. They manifested themselves as emergent properties at the end of the design and development phase, requiring additional design interventions and operational remedies (Stoop and Dekker, 2007).

- The institutional environment has complicated the development and implementation of the project. The divisions that were created during the project between design and construct of the hardware components and the contractual arrangements between stakeholders required a complex interface management. This interfacing has not been accomplished.
- The necessity to create oversight emerged only by the end of the project. There was no role for a systems integrator, responsible for the integral coherence of the overall system. The pivotal role of ERTMS became emergent at the end of the project in the full scale testing phase of the integral system.
- The technological development of ERTMS was underestimated. There has been a continuous tension between incremental progress and implementation in an existing railway network on one hand and the ambitions of innovative ERTMS and public-private partnership arrangements on the other hand.
- Consequences of several technological design decisions should have been submitted to a pro-active safety assessment procedure. Several points-of-no-return in the design process were passed without oversight of their consequences.
- A choice for a new signalling system, which was not yet operational at the time, was not compensated by a qualified fall-back option.

- The choice for an innovative ERTMS system in the Netherlands was not in harmony with the more incremental process and evolutionary development of the Belgian signalling system on the same corridor Amsterdam-Antwerp.
- The choice for connecting the Dutch and Belgian system manufactured by two different signalling system consortia at the country border forced the project management to develop a gateway causing high costs and considerable delays in delivering the integrated system for testing and operations.
- A contractually based testing and deployment of ERTMS version 2.2.2 took place while version 2.3.0 would become the new standard, causing complications, costs and delays.
- The development of ERTMS was considered a conventional technical engineering effort, enabled by a decomposition of the system components in autonomous position finding and communication subsystems. Development and manufacturing of these components was subcontracted across competitive consortia. Each consortium was assumed to be able to deliver these components 'off-the-shelf' as proven technology.
- No precautionary measures were taken to assure a smooth and efficient frequent upgrade of the signalling software during its operational phase.

Transparency

Safety can be considered a normal consequence of performance variability which should be controlled rather than constrained (Hollnagel et al., 2008). The ability effectively to adjust a system's functioning *prior to* or *following* changes and disturbances can sustain its functioning after disruption or mishap, while under continuous stresses. Systems should therefore be able to cope with responses to the actual, critical, potential and factual situations. A transparency in various system states should be available supported by the ability to predict, plan and produce.

But what if such transparency is impossible? If we cannot analyse the complex reality and cannot achieve consensus, are we doomed to restrict ourselves to a battlefield of subjective opinions, submitted to political will and governance resolve (Rosenthal, 1999). Or do we restrict ourselves to a lower level of a single agent at the organisational level, accepting sacrificial losses? The potential of offering opportunities in solving complex problems by taking into account the dynamics, multidisciplinarity and complexity of systems enables a transition from a static and event oriented approach into a dynamic, system oriented approach by applying chaos and complexity theory and a re-introduction of the conceptual design phase in system change (Bertuglia and Vaio, 2005). This dynamic system perspective identifies systems behaviour beyond the level of linear behaviour dealing with deterministic chaos, emergent behaviour, self-organisation, self-conformity, resonance, and bifurcations. From a safety perspective, the most interesting parameter is the existence of multiple system states (Hendriksen, 2008). This perspective eliminates the debate on acceptable and quantifiable system safety performance levels, replacing it by the need to design resilience into such systems in order to cope with change.

Towards a New Train Control Concept

The Resilience Engineering potential has been demonstrated in a feasibility study into the deployment of a new railway concept beyond the boundaries of present railway configurations. It aims at doubling the number of trains for half the costs per passenger kilometre, maintaining present safety performance levels. Regarding the train control functionality, a new Free Ride concept was developed in analogy with the Free Flight concept in aviation (Van Binsbergen et al., 2008). Instead of the full automation paradigm, a human-centred design approach is applied.

The Free Ride concept eliminates the conflict of interest between safety and control, by applying a performance based local control strategy instead of a centralised compliance based approach, restricting incident management and handling to the local level of the network (Van Binsbergen et al., 2008). Such a concept allocates the incident handling responsibilities and control options to the local level, which only can be overruled by a centralised control to avoid a network gridlock state. Four innovations of both a technological and organisational nature are required for a further development of the Free Ride concept (Van Dam et al., 2008; Van Eijndhoven, 2008).

Lessons Learned

In comparing the safety assessment of the ERTMS system development against the Free Ride concept, some conclusions can be drawn.

- Actors located at different systems levels and life cycle phases create value and control conflicts. A multi-agent process approach is emergent, but not sufficient to guarantee a continuous and explicit interest in safety issues. Safety is a system critical aspect that cannot be reduced to a field test item in a final phase of the project.
- Technological innovation creates major uncertainties. Engineering is not a standard technology application which can be bought off-the-shelf; it also contains software design concepts change, system architecture and integration, oversight/consequence analysis and integral system certification and testing.
- Shifting from a technological perspective in systems development towards a social engineering is not sufficient; there is a need to integrate the technical, human and organisational/institutional design across the various system life phases, taking into account the various system states that may exist in practice.

Discussion

Engineering design aims at expanding the design scope from form to function and from performance to properties, from aircraft design towards systems design, and from component towards context.

Taking this integrated engineering design approach facilitates identification from emergent properties in practice to inherent properties in the design phase. Integrated design is dealing with dynamic instabilities in the overall Programme of Requirements.

As such, resilience is a property of saturated and mature systems, emerging during the operational phase. It is a major challenge to deal with this system property already in the design and development phase.

What Do We Need to Design Resilient Systems?

In order to be able to design resilience into systems, several requirements have to be met.

- Unravel complexity in the system by investigating system levels, life cycle phases and safety critical decisions, components and interdependencies in order to provide transparency in the systems functioning as designed and as operated and identify how hazards may propagate through this system. Triggering events may provide opportunities for change by creating a sense of urgency, but do not necessarily focus on critical deficiencies.
- Identify system and knowledge deficiencies in order to understand and control the propagation of hazards through a system. Such investigations should be unbiased and impartial. Instead, an understanding of goals and motives should facilitate a perspective on improved systems governance and control and should facilitate change strategies.
- Apply an integral systems approach from a multi-actor and multi-aspect perspective in order to achieve consensus across stakeholders in a common understanding and ability to change the overall system's performance. Such a perspective reinforces the resilience of a system against unanticipated changes in internal and external operating conditions and constraints.

- Acknowledge the specific role of technology as a mature and often saturated system characteristic that is submitted to a need for innovation at a conceptual level in order to make substantial progress in performance. The need for a system oversight requires an explicit role for an integrator or architect for the system. In particular in aviation, technology has been selected as the flywheel for progress (Freer, 1994).
- Such a need for innovation and systems integrator role cannot be restricted to the technological components: organisational and managerial changes and institutional arrangements have to be taken into account as well. Expanding the design solution space refers to the overall systems design level and integration of new functions and actors.

What has Created Opportunities for Resilience in these two Cases?

In the Schiphol case as well as in the ERTMS case resilience has been designed into the system by:

- Expanding the design solution space beyond existing boundaries. The role of innovation in a technological, organisational and institutional sense has been crucial for achieving new integrated solutions. A new technological concept for wide body aircraft, airport infrastructure, train signalling and control had to provide the necessary technical capacity for the intended growth and expansion.
- Recognition of the need for systems oversight and system integration. Optimising components and single functions does not fulfil the need for an optimum overall performance of such complex systems. Public values such as safety may easily be jeopardised and sacrificed against corporate and private values such as economy.

- The familiarity with a broad range of disciplines and the ability to pursue several lines of investigation simultaneously allows investigators and designers to apply multiple perspectives and to ensure democratic participation in achieving consensus on goals, strategies and assessment of the final results.

In retrospect, the investigations into the air crash and the railway signalling system both provided transparency in the systems, revealing deficiencies which were unnoticed until the investigations were conducted. As such, the investigations offered additional options which came available for enhancing the systems safety. Without these facts, there would have been no glory.

This page has been left blank intentionally

Chapter 16
From Myopic Coordination to Resilience in Socio-technical Systems. A Case Study in a Hospital

Anne Sophie Nyssen

In socio-technical systems, the overall functioning requires, by definition, a coordination of actions and decisions of the different agents involved in the task. It seems reasonable to assume that the resilience capacity in such systems therefore also should require the coordination of local and spontaneous coping strategies distributed in time and space to recover from surprise, unexpected events or crises. Indeed, the history of accident investigations contains many instances of dramatic coordination failures. The purpose of this chapter is to show why the study of coordination mechanisms is so crucial to the resilience of socio-technical systems. We illustrate this importance using an example from one kind of socio-technical system, a health care system. First, we introduce the concept of coordination as a component of resilience in socio-technical systems like hospitals and show how they handle coordination requirements. Then, we describe how practitioners adapt to the unexpected using emergent coordination mechanisms. We conclude by developing our main argument that resilience of socio-technical systems largely depends on their ability to project themselves outside the 'local immediate' when an unexpected event arises in order to be able to develop a coordinated spatio-temporal solution.

Introduction

In socio-technical systems, the overall functioning by definition requires a coordination of actions and decisions of the different agents involved in the task. It seems reasonable to assume that the resilience capacity in such systems therefore also should require the coordination of local and spontaneous coping strategies distributed in time and space to recover from surprise, unexpected events or crises. Indeed, the history of accident investigations contains many instances of dramatic coordination failures. All organisations, including socio-technical systems develop capabilities to detect and deal with unexpected events. These are part of the uncertainty of the world that every complex system must face and learn to cope with. Part of the capabilities developed by such systems includes the development of rules, procedures, standardisation and training programmes. These centralised systems organise and control the interactions both inside the system and between the system and its environment in order to maintain continuing alertness and to control the safety boundaries of the organisation. However, it is accepted that it is impossible to anticipate and write down a rule for every circumstance. Furthermore, the collection and analysis of unexpected past events, such as accidents, do not allow the prediction of future ones.

In that context, system resilience also seems to depend on the front-line agents' capacity for flexibility, local autonomy and creativity, which allows them to adapt to changing and unexpected circumstances. This flexibility is generally considered as a positive contributor to – or even the base of – the resilience of complex systems, through the 'loose coupling' they introduce between the system's agents and the internal and external constraints. However, contrary to frequent assumptions about sharp end performance, local, spontaneous actions and decisions are not always virtuous (positive).

In our view, resilience in socio-technical systems relies more on the dynamic of the interactions between the agents than on each individual or sub-component's ability to adapt. It is the dynamic of these interactions that allows such a system to cope successfully (or not) with unexpected events, irregular variations or crises.

The purpose of this chapter is to show why the study of coordination mechanisms is so crucial to the resilience of socio-technical systems. We illustrate this importance using an example from one kind of socio-technical system, a health care system. First, we introduce the concept of coordination as a component of resilience in socio-technical systems like hospitals and show how they handle coordination requirements. Then, we describe how practitioners adapt to the unexpected using emergent coordination mechanisms. We then examine some conventional dimensions of safety and resilience and show how they may fail in socio-technical systems. We conclude by summarising the main arguments.

Coordination as a Component of Resilience in Socio-Technical Systems like Hospitals

Coordination in socio-technical systems covers two parts: a movement of division and distribution of actions among different agents and a movement of integration of actions and decisions distributed in time and space (Savoyant and Leplat, 1983, Pavard, 1994). All socio-technical systems, like hospitals, are confronted with both. Today, a patient will very seldom visit only one hospital department and furthermore rarely sees only one physician during their stay. Multiple departments, professional skills and technical devices are brought together in order to provide a complete health service and also to provide uninterrupted care around the clock. This specialisation and continuing process means that more and more information has to be exchanged between departments as well as between individual operators. Hospitals themselves have even become specialised because of economic pressures, so that a patient may have to go to several hospitals and health care institutions to be properly taken care of. This obviously raises the coordination challenge to the inter-organisation level and makes it necessary to also consider resilience at that level.

Classically, we can distinguish the situations of coordination on the basis of four dimensions (Bonabeau and Theraulaz, 1994):

- the compatibility of the goals of each agent involved in the task
- the sharing of common resources
- the skills of the agent in relation to the task
- the type of interaction – face to face and synchronous communication (same place and same period) or distant and asynchronous communication (different place and different period).

Each may be a source of tension in socio-technical systems when an unexpected event arises and hence have implications for resilience. For example, confronted with the same patient situation, the surgeon would prefer to operate as soon as possible while the anaesthetist would prefer to stabilise some parameters. There is one common goal related to the patient's health but two contradictory sub-goals related to the emergency strategy. Conflicts may also appear when all the agents actively decide to cooperate, for example, when different practitioners are called to an operating room to deal with a cardiac arrest. Their actions and decisions must be coordinated if we don't want this assistance to end in chaos. The agent's experience in relation to the task also influences the exchange of information and the coordination. When practitioners repeatedly work together, a reduction of verbal information exchanges is observed as practitioners get to know each other. Any regularly repeated action by a member of a team becomes an act of implicit communication, a signal that triggers actions by other team members, hence allowing synchronization, just as explicit verbal communication would do (Nyssen and Javaux, 1996). This relation changes during an unexpected event or crisis: the greater the trouble, the greater are the demands for information centred on the task (Bressole et al., 1994).

Another major change for communication and cooperation mode in hospitals comes from the introduction of new computer-based technology that allows distance between the operators and the task (e.g., robotic surgery) and/or between the operators (e.g., electronic patient's file). These technologies deeply transform the coordination situations from face-to-face and synchronous communication to distant and asynchronous communication in which information is mainly built up incrementally by accretion. In robotic surgery, we have showed how the technical device changes profoundly the structure of the surgeons' task and, hence, the mode of cooperation between the surgeon and his assistant. It favours an explicit division of work and an explicit leadership based on order communication and continuous control of the work or asks for confirmation. Whether the interaction is synchronic or asynchronic may be a critical factor in view of the emergency response capacity of an organisation (Johansson and Hollnagel, 2007). When confronted with unusual problems, people in charge are naturally prompted to intervene, although the problem could have required sharing information distributed over time and space and a collective response.

These developments help us to envision the coordination challenge for socio-technical systems confronted to unexpected event and crisis. But, they do not specify how organisations deal with these challenges.

The Organisation's Approach to Coordination

When we consider the organisational side of coordination, we can identify three sets of coordination requirements: vertical, lateral and longitudinal.

- *Vertical coordination requirement*. All organisations, including hospitals, structure their work and activities through a vertical distribution of decision-making responsibilities and roles. Because of their various positions in the hierarchy, individuals are likely to have differential access to information and knowledge. The exchange of information is required to keep the subordinates' supervisor aware of the situation and decisions. A vertical flow of information is needed, both bottom-up and top-down, to provide the right people with the right information to carry out the task.

- *Lateral coordination requirement*. Another important aspect of complex organisations is the existence of different fields of expertise that are institutionalised and organised into different units or departments, which have with different technologies and specific subcultures. In the collective activity, each expert has only a partial representation of the situation. The task requires the coordination of the different expert's activities and, in many cases, the evaluation and integration of the information from these various sources into one global base of knowledge, if not a shared representation.
- *Longitudinal coordination requirement*. Many organisations operate all around the clock, requiring exchange of information between the different skills. The tasks themselves are made up of subsets of activities or sequences of actions that must be executed in the proper form and with the appropriate timing. These process constraints shape the coordination of work either at the longitudinal or the lateral coordination levels. Furthermore, many processes are dynamic and are subject to modification; the supervisor must continuously update his/her representation of the situation.

Any organisation positively organises its coordination by developing a series of conventional management tools that specify *ad hoc* patterns of behaviour and directly or indirectly shape the interactions and the communication between the agents. These tools supposedly enable the agents to manage everyday and unexpected situations more effectively.

This coordination approach is based on the Common Knowledge Theory (Lewis, 1969, Krauss and Fussell, 1990) and is derived from the classical assumption that the success of coordination lies in the extent to which the community and individual agents are prepared to understand and share 'common ground'. In work studies, the idea of 'common ground' shared by the individuals who perform a collective activity led to the concepts of 'functional referential' (Leplat, 1991), or 'mental model' (Norman, 1987) that define work processes and allow group members to organise their activities.

For example, the aviation industry has attempted to reduce the problems of cooperation between humans and automatic systems by organising both human–machine and human–human communication, using a straightforward and predefined division and distribution of tasks (e.g., Pilot Flying and Pilot Non-Flying), a codification and standardisation of the communication language, a principle of systematic verbalisation of the main intentions, perceptions and actions (call-outs), a principle of systematic cross-checking of actions and understandings, and mandatory training of so-called 'non technical skills' (Crew Resource Management).

Although health care practitioners point out that 'improving communication' is an important corrective strategy (Kluger et al., 2000), communication has not received much attention in hospitals. Better training, better techniques and better standards of equipment have been recommended in order to improve the patients' safety, but not much effort has been made on communication training and tools. Coordination then relies more on a series of conventional management tools such as hierarchy, work organisation processes, patient processes, procedures, daily lab rounds, patients' files, handover meetings and the like. We can identify such conventional tools for each coordination level.

Vertical Coordination Tools

The work organisation in a hospital allocates the simplest part of the collective patient task to the novices and the more complex part to the more experienced workers. Novices' work is commonly monitored by residents and/or seniors. In many hospitals, a phone communication network is organised in a cascade to provide, at all times, help from those who are more experienced. The basic functioning of vertical coordination is that the novices do their work, with some internal and external 'sentry markers' (events, parameters) which tell them that it is time to call for assistance from someone at a higher level of expertise. The key issue is therefore the relevance of these 'sentry markers' that alert the novices and suggest a call at the right time. A previous study (De Keyser and Nyssen, 1993), showed that trainees often fail to estimate correctly how long they should reasonably wait before calling and do so too late: either they overestimate their competences or they underestimate the speed of the patient's deterioration or sometimes they do not want to lose face.

Lateral Coordination Tools

The resources of the hospital (either technical or human) are both specialised and limited and a careful coordination of the activities, both in time and in space across the facilities, is required. This resource management is based on multiple planning activities, which are organised according to different time frames: they may start up to a year beforehand, and then evolve from annual, to monthly, to daily and hourly schedules. Computerised systems are used to exchange information between people from different areas in order to gather the right people at the right time and place and achieve lateral coordination. The work process itself also defines how the tasks are organised among the agents and shape their interactions. Individual and collective patterns and sequences of behaviour are defined into procedures.

Longitudinal Coordination Tools

There are rotations of multiple teams in charge of the patient around the clock in a hospital. A transition period is generally planned in the agent's schedule to allow for data transmission briefings.

One important tool for longitudinal (and lateral) coordination is the patient's records, either in its paper form or its computerised form. For both the hospital and for the team, it is a means to trace and memorise the state of the patient and the actions of the different agents around the clock. It contains the history of the case, contextual information, the distributed dynamic diagnosis and the treatment process. Each agent is supposed to fill in the patient's records with their contribution to the actions and information and to transmit their knowledge to the next agent. By accumulating information over time and from different agents, the patient's records play a critical role in coordinating. It is intended to produce the global representation of the patient's situation to enable an isolated agent to solve dynamic problem situations.

The hospital normally takes for granted that the staff will adhere to these coordination principles and, so doing, assumes that coordination problems are solved. However, in the following case description, I will show how these centralised coordination mechanisms fail to organise the activities of different agents when unexpected event arises and, hence reveal the adaptation strategies and the resilience capacity of the system.

A Catastrophic Experience

One night on a weekend, a 16 year old patient showed up at the reception desk of the emergency room of Hospital A for respiratory distress related to a problem of chronic asthma. The clinical examination was done by a resident who prescribed treatment (inhaled bronchodilator and antibiotic therapy) and let the patient go home.

Later in the night, his respiratory distress symptoms reappeared at a higher degree. In the early morning, the patient showed up at the emergency room of a larger hospital (B) in the same area. The patient was in an agitated and anxious state. He was directed to a room of unit X where he was examined by resident 1 (R1) and monitored (ECG, arterial blood, pulse). The clinical examination did not show any acute respiratory problems. R1 gave the patient some treatment and kept him for close monitoring. At the end of his shift (9 am), R1 transmitted all the information concerning all the patients to the resident 2 (R2). In the middle of the morning, the patient felt better and R2, after examination, decided to let the patient go.

In the early afternoon, the patient showed up again at the emergency department of hospital B, still complaining about respiratory distress. At the reception desk, the secretary recognised the patient. For her, there was no real emergency and this time she decided to refer the patient to Unit Z. Unit Z had just been created in the emergency department in order to take care of minor emergencies. General practitioners (GPs) from outside the hospital were used in that unit for patient care. In that unit, the patient was examined by a GP who was on call in the hospital for her fifth time. The nurse, who was working with the GP, also recognised the patient and informed her about the case. The nurse went to get the patient's records and, at the same time, asked R2 to come and see the patient.

The GP and R2 were both in the room and examined the patient. The patient showed some signs of tachycardia, nasal flaring, hypoxemia and anxiety. R2 decided to give the patient oxygen and started corticoids and the GP proposed to give him some anxiolytics. In the afternoon, the patient complained about chest pains. The GP added some analgesic to the treatment and called a psychologist who did not diagnose any particular mental problems.

At the end of the afternoon, a nurse came to see the GP to tell her that the patient was not complaining anymore and seemed better. She asked her to come and examine the patient in order to see if he could go home. After close examination, she decided to let the patient go. R2 heard later that the patient had left.

In the evening, the patient went into respiratory arrest at home. He was taken by ambulance to hospital B but too late.

Case Analysis: An Emergence-through-use Approach of Coordination

As for many accident analyses in complex systems, it is not easy to identify where and how the case went wrong through the course of events. Leaving aside the benefit of hindsight, each decision, each action seems to be relevant in its finite temporal and spatial interval. The overall failure appears to come from a lack of integration of the decisions and actions distributed in time and space and across the agents; that is to say from a lack of coordination.

The centralised organisation of coordination presented above failed to achieve the level of integration of knowledge and action that was necessary for handling the problem situation. Confronted with an unusual situation, the agents organised their behaviour through direct and local interactions with the work environment, based on their understanding of the situation. However, this local process of coordination disorganised the standardised sequence of operations. Let us analyse in detail the dynamic of the interaction and the coordination failure.

- Hospitals A and B are two units in the same urban area. Both are capable of taking care of a chronic asthmatic patient. For such a 'routine' problem, they are two equivalent resources, two interchangeable structures. Hence going to hospital A first, then to hospital B should not have been to be an issue for the patient. However, coordination is not organised between the two units. There was no communication of any type (verbal or written) between the two hospitals Consequently, each unit re-starts the diagnosis process: creating its own patient records, its own diagnosis and its own treatment process. Each diagnosis process seems to be relevant in its finite temporal and spatial interval. But because of the slow dynamic of the problem, the move from one hospital to another impaired the detection of the overall representation of the problem. However, it must be noted that this lack of communication could have been overcome by the application of a recent national measure that appoints the patient as the agent who keeps his exam record files.
- There is a procedure written by the chief of the department that organises the patient orientation across the different sub-units of the emergency department, and so lateral coordination. However, this algorithm was not used by the agents in the case described above. They were not really informed about it, and they were not involved in its development process. Furthermore, the algorithm does not cover the case of a patient who comes back again several times in the same day. Actually, the first line agent, the receptionist, achieves coordination without formal procedure. The receptionist identifies the degree of emergency, the nature of the problem and attempts to match the demands with the available resources. This

matching is mediated by direct interactions with doctors and nurses, and supported by a computerised system that gives an external representation of the workload of the different sub-units. This global representation of the team's activities is permanently updated, yet does not keep track of past activities. Its goal is to help the synchronic management of resources. It was not intended to help the agents for longitudinal coordination.

- The coordination between GP and R2 is something that emerges from a set of local interactions rather than by the implementation of the procedure that explicitly organises the transmission of information between physicians. In the case study, it was the nurse who detected the presence of the patient the second time and organised the transmission of information through the patient's records and by direct verbal interaction between R2 and GP.

  By doing this, the nurse created a system in which everyone believed that the others knew everything about the task, hence shared the same understanding in a fully cooperative work. This emergent movement of increasing expertise resulted in a pattern of overlapping expertise rather than of cooperating work. In fact, the knowledge of the medical task was represented most redundantly, but the centralised knowledge necessary for a dynamic problem solving was paradoxically represented least redundantly, and was lost across the agents as well as in vertical and lateral coordination.
- Each physician used the patient's records but everything happened as if each agent started a new reasoning process instead of integrating the information recorded in the patient's records and constructing a global dynamic representation. Clearly, the individual performance of the agents was not improved by the use of this external static memory source assumed to achieve longitudinal coordination. Even the direct verbal interaction among individuals did not really help to alert the physicians and provide any diagnostic benefit. In contrast, these local interaction and local coordination processes might

> actually have worked against the efficiency of the problem solving process by confusing each person's role and creating some kind of 'stammering' in the reasoning and treatment process.

Discussion

Part of the benefit of being a socio-technical system facing a crisis or an unexpected event comes from the juxtaposition of agents, either human or technical. The different agents create redundancy; each agent can detect signals or dangers, update the process representation with new information and interactions with the environment and formulate a regulation plan. The increase in the number of agents has by its nature an adaptive value for the system, at least to a certain extent. This is a very simple form of redundancy relying on the increase in resources for signal detection, diagnosis and action plans, bringing benefit in terms of resilience. But, there is no increase in expertise among the agents who are interchangeable. In our case study the patient himself, when confronted with an unexpected evolution of his symptoms, adopted this strategy by increasing the number of hospitals he interacted with. A beneficial aspect of this strategy is that by replaying the game every time, each agent can detect someone else's error and improve safety. However, our case study shows that this strategy may not be optimal in some problem situations.

The delivery of appropriate medical care depends on obtaining information from different sources that could specify the cause of the patient's symptoms. In many cases, this is an iterative task. The complexity comes from the fact that these sources are distributed in space and in time. When our patient decided to move from one hospital to another, he involuntarily created a rupture of this task, affecting the critical integration of the temporal aspects of the problem. The detection of the dynamic pattern of the problem and its repetition over time was impaired by the distribution in space and in time of the diagnosis and decision making process. The health care system lost its resilience.

A second approach relating to resilience may be found in the way hospitals have dealt with the coordination problems by developing centralised tools such as written procedures, work processes, automated systems that specify the work and the activities across time and distance and guide interactions. The above situation shows how these centralised coordination mechanisms may fail when the tool does not cover a particular case. Reasons for these failures can be, for instance, that the agents are not familiar with the conventional tools, that the computerised systems are not designed with the coordination needs in mind and so the critical information for coordination is not saved or not transmitted.

Correlatively, our case study clearly demonstrates that when agents are faced with an unexpected event, they often rely on 'emergence-through use' coordination instead of referring to centralised tools. Each agent seems to organise the activities through direct and local interactions in his/her work environment.

By analogy, recent research that studies coordination in insects and non human societies shows local coordination mechanisms at the origin of very complex patterns of adaptation (Bonabeau and Theraulaz, 1994, Gilbert and Conte, 1995). For instance, Reynolds (1987) demonstrated that the flocking behaviour of birds can be simulated by assuming that individual birds make local adjustments based on the velocity and bearing of neighbouring birds. Thus, despite appearances, such complex flocking behaviour is generated by local coordination processes rather than by global centralised ones.

A strong message from High Reliability Organisations' commitment to resilience is about their sensitivity to the front end operations and their ability to distribute the control and decision making to the low level members (Weick and Sutcliffe, 2001). These members are closest to the problem and are better able to adapt as the tempo of operation quickly changes and unexpected problems arise. A central assumption is that when people have a well developed situational awareness, they can make the continuous adjustments necessary to respond to the dynamics of the situation and the unexpected. This flexibility of the decision structures is an important issue for resilience in large scale organisations like hospitals facing the unexpected.

But local regulations are not always positive, contrary to some naive assumptions about sharp end performance. In robotic systems, Mataric (1992) has shown that distributed control can lead to a kind of 'myopia' (short-sightedness). When there is no centralised supervisor that possesses a global representation of the operations, the group of robots can fall into the trap corresponding to a 'local minimum.' We have shown in our case study that each agent seems to give precedence to their own current perception of the situation based on his/her local and real time interactions with the patient, and re-starts the reasoning process instead of continuing it, falling into the trap of 'myopia'. Within each spatio-temporal window, the problem appears as a 'routine' emergency problem and each physician copes adequately with the emergency symptoms. In emergency departments, people in charge naturally focus on emergency but in doing so the system fails to capture the global pattern of the problem. Apparently, the patient was well at the hospital, went home and the symptoms started again. The process is not linear. The challenge for dealing with the problem is to understand this pattern. Otherwise, the management of the problem will always stop too early. This understanding would have required that, beside the emergency symptoms, one physician sat back, shared questions with the other physicians in charge, reconstituted historically the different isolated responses, integrated the external factors' influence and anticipated the lethal process when the patient is pushed outside the hospital. However, as Lagadec (2004) mentioned in his analysis of the 2003 French heatwaves that killed nearly 15,000 people, in most cases, the culture during crisis is you act – you do not have time to think.

Conclusions – 'Enhancing Projection outside the Local Immediate'

From a development perspective, there is a difference between everyday life coping strategies and resilience capacity. The difference is not the intensity of the event but its pathologic impact that requires a reorganisation of the system in order to maintain safety and survive the crisis (Bowlby, 1973, Cyrulnik, 2003).

The issue in this chapter is not to argue for or against one of the two organisation's approaches of coordination described above: centralised or emergent. It is evident that the two approaches of coordination are clearly embedded in work practices and can both be beneficial in terms of resilience. However, these approaches may not be sufficient in today's large scale organisations and should be complemented by the capacity to coordinate responses over time and space.

My argument is that resilience of socio-technical systems largely depends on their ability to bypass the myopic cognitive bias mentioned above in order to be able to develop a coordinated spatio-temporal solution. This requires two competencies from the systems: the ability to cope with the unexpected as it arises (this was effectively done in our case study partly thanks to the local spontaneous coordination processes), and the ability to keep on projecting themselves into the future beyond the present but taking into account the past.

Projection refers to a process of symbolisation of the diversity and complexity of all eventualities. It requires the ability to keep interactions going internally as well as with the external environment during the crises, in order to be able to capture the history, to read the changing circumstances and to be prepared for what the future holds. This is fundamental for coordinating the responses over time.

Following this argument, a rich array for future research for resilience design is a better understanding of how to enhance systems to project themselves outside the 'local immediate' when an unexpected event arises and thus how to represent and record the histories of the local agent-environment coupling adjustments distributed in time and in space in order to enhance a coordinated spatio-temporal regulation process over time. What kind of tools are the most appropriate to tackle this new coordination requirement? We have seen that traditional tools allow data saving and sharing, but mainly in an asynchronic mode. In our view, this discontinuity may favour the repetition of the regulation process instead of its steps-by-steps refinement over time. The challenge is to inject learning into the adjustment process. This may require synchronic interaction for collective decision-making expertise to elaborate coordinated responses over time.

From this perspective, the study of interaction becomes an important paradigm to capture the resilience capacity of socio-technical systems. The idea of interaction as an instrument of development of cognition, and thus serving adaptation is not new. It is central to Piaget's theories (1967, 1992). Adaptation, in his constructivism framework, is achieved through agent-environment interactions via the conjunction of two processes: (a) the assimilation of new experiences into existing structures, and (b) the accommodation of these structures, that is, adaptation of existing ones and/or the creation of new ones. The latter, learning through accommodation, occurs for the purpose of 'conceptual equilibration' and the elimination of perturbations.

At a metaphorical level, the resilience capacity in socio-technical systems becomes observable and defined through the study of interactions and coordination modes inside the system and between the system and its environment.

This page has been left blank intentionally

Chapter 17
Requisites for Successful Incident Reporting in Resilient Organisations

Alberto Pasquini, Simone Pozzi, Luca Save and Mark-Alexander Sujan

This contribution offers a critical reflection on standard reactive incident reporting systems and provides an outlook towards proactive methods for monitoring risk. Incident reporting systems are often regarded as a prerequisite for effective Resilience Engineering, but sometimes they fail to achieve most of the expected benefits. There is now a growing body of research that criticises incident reporting on the basis of its inability to provide an accurate representation of harm compared to other methods, as well as the fact that there is still widespread under-reporting of incidents. In this chapter we take a different angle by arguing that the problems encountered with incident reporting are, at least to some extent, to be found in the structural characteristics of the respective domains rather than within either the principle of incident reporting as such or its implementation. We identify a number of such structural characteristics that are necessary for successful incident reporting through reflection on the success and (partial) failure of two major incident reporting systems from aviation and healthcare. Where those structural characteristics are not present, incident reporting systems are bound to encounter difficulties. In such environments, a complementary proactive risk monitoring approach may be required to maximise learning from operator and front-line feedback.

Introduction

A systematic approach to safety management has greatly improved the safety performance of many safety-critical systems, to the point that very few serious accidents happen in domains like railways, aviation or nuclear processes. However, such systems face a contradiction inherent in their excellent safety performance: How can we continue to learn from accidents if we succeed to prevent most of them? The paradox is that a zero-accident system loses a valuable information source by improving its safety performance and it needs to replace it with some alternative source. One of the well recognised solutions to this contradiction is the establishment of an incident reporting system. Incident reporting systems have been devised to make sure that continuous learning is in place by relying on operators' feedback (Johnson, 2003; Reason, 1997; Van der Shaaf et al., 1991). Operators are in the best position to closely monitor system performances and to detect any deviation from normal operating conditions. They can be asked to report all the near-misses, that is all those cases when an accident could have occurred but was avoided by operators' intervention, or even by fortuitous circumstances. This is especially true with respect to system evolution, in the sense that operators do not only recognise existing unknown hazards, but they can also closely monitor how the system changes, due to external or internal forces.

The most prominent experience in incident reporting is the Aviation Safety Reporting System (ASRS), established in 1975 by the Federal Aviation Administration (FAA) and the National Aeronautics and Space Administration (NASA). This is often cited as best practice in incident reporting, due to its lengthy duration and to the fact that it is almost unanimously regarded as a useful system. Unfortunately, even though many other attempts have been made in various domains to establish similar reporting systems, very few are able to claim a similar success. Different success and failure factors have been discussed in the literature, spanning all levels of analysis and explanation, for example, usability of reporting forms, organisational structure, national legislation, operators' lack of involvement, etc. (Johnson, 2002). Many efforts aimed at improving the performance of incident reporting systems are therefore directed at the way incident reporting is implemented (user-friendly forms, user involvement during the design, etc.) and the cultural environment within which it is implemented (open, fair and just culture, feedback to reporters, etc.).

This chapter discusses the introduction of reporting systems as a socio-technical issue, that is, by considering reporting systems as embedded in safety critical systems and domains, whose characteristics have an effect on the reporting system efficacy. Depending on the nature of the system or domain, a different approach to reporting and utilising operator feedback may be necessary in order to maximise an organisation's capability of learning from experience. Such an approach – risk monitoring – proactively elicits feedback from operators about the dynamics of variation and risk present in the system.

A Success and a Failure Story: Reporting Systems in Aviation and Healthcare

This section analyses two reporting systems in two different domains. We first review what is currently referred to as the best practice in incident reporting (i.e., the FAA reporting system) to reflect on why this system is capable of collecting good quality data and of transforming them in actionable recommendations. We then compare this case with the UK National Reporting and Learning System (NRLS) introduced in the healthcare to improve patient safety learning.

The Aviation Safety Reporting System

The Aviation Safety Reporting System (ASRS) is an independent system, run completely outside the FAA. Its main objectives are to:

- discover patterns of frequent problems;
- improve communication on major issues;
- support policy making with empirical data.

Pilots, air traffic controllers, flight attendants, mechanics, ground personnel and others involved in aviation operations submit reports to the ASRS when they are involved in or observe an incident or situation in which aviation safety was compromised. All submissions are voluntary. Reports sent to the ASRS are strictly confidential.

The core part of an ASRS report is a narrative of the event. This is provided in a free text format. Other fields to be filled in contain more standardised information, like for instance the airspace type where the event occurred, the phase of flight, date, time, geographical location, etc. Each report thus contains factual information about the event – where factual does not mean 'objective information', but rather descriptive information about the event with no further elaboration, for example, no causal factor analysis. The free text format indicates that no strict instruction is given on what should be included. The reporter is supposed to write everything they think is appropriate, including all the required details. These narratives provide an exceptionally rich source of information for policy development and human factors research.

The process of analysing the reports can be divided into two steps. First, all reports received by ASRS are reviewed by two analysts. Analysts are experienced pilots or air traffic controllers. Each report is screened against established criteria to determine if it warrants full analysis and if it should be entered into the database. Currently, 25–30 per cent of reports pass this screening and are inserted into the database. Reports that undergo full processing fall into four categories: (i) aviation hazards that require immediate alert messaging; (ii) priority safety concerns that have been targeted for data collection; (iii) random sample to ensure database representativeness; and (iv) reports that, based on the discretion of the expert analyst, represent a new or unique learning opportunity.

After the initial screening, the report is further analysed. The first aim of the analysis is to identify any aviation hazards and flag that information for immediate action. When such hazards are identified, an alerting message is issued to the appropriate FAA office or aviation authority. The analysts' second mission is to index reports and diagnose the causes underlying each reported event. An important point to mention here is that this may also imply that people involved in the event are contacted to gather further details or clarify key points, as one of the goals of the analyst is to make sure that the narrative is descriptive, complete and precise. The system is thus confidential, but not anonymous and reporter identity is discarded only after the analysis phase has been closed.

In the previous description of the ASRS, we have hinted upon structural domain characteristics that are of paramount importance for the success of the ASRS system. In other words, behind the successful achievement of the main ASRS objective (gathering operators' point of view to discover unknown system weaknesses) we should not downplay the role of particular domain characteristics. Three of these appear more relevant.

First, in the aviation domain there is a clear-cut distinction between an incident and an accident, and between incidents and non-relevant events. Only incidents should be reported to ASRS, while accidents are investigated by the legally entitled authorities. In a similar way, operators know how to distinguish mundane disturbances from real system weaknesses. This is indicated by the fact that even if up to 75 per cent of the reports are not warranted to need full processing, still the system is not overflowed by irrelevant reports. To oversimplify the point for clarity's sake, in the aviation community there is a shared agreement on what is a safety relevant fact and on the criteria to assess its severity (this characteristic will be later referred to as the 'pass criterion').

Second, well defined roles and professional communities are present in the aviation world, meaning that the ASRS can put together a complete team of experts to represent all the different points of view. The ASRS is considered both as an independent external organisation and as possessing the relevant expertise to conduct the analysis. ASRS analyses are seen as trustworthy and competent by the aviation community, which implies that the community is to a certain degree open to an 'external' judgement as long it comes from a recognised expertise. This trade-off between being independent and external, but still preserving the required expertise, is often encountered in safety critical domains, for instance in cases of investigations, of safety relevant data gathering, of regulatory bodies, etc. The solution is often very hard to achieve as a body that is too independent can fail to be recognised as competent by highly specialised professional communities, while a institution that is too internal tends to reason too similarly to the community it should oversee (this characteristic will be later elaborated under the heading 'Understand the characteristics of your community').

Third, given the high degree of standardisation of aviation operations, textual narratives are considered a good means to describe the event and to conduct the analysis. The high level of standardisation ensures that contextual factors can be omitted in the description, as the analysts will be able to fill in for themselves this background information without the reporter explicitly describing it. This also implies that the community knows to a reasonable extent what are to be considered 'normal operating conditions' and what should be regarded as a non-standard event deserving full description (cf., the discussion of 'Degree of standardisation' later).

Incident Reporting in Healthcare

In this section we will focus in more detail on reporting systems in healthcare, with a particular focus on the UK National Reporting and Learning System (NRLS), the only national system currently in existence (comparable, though not truly national systems, include the Veteran Affairs system in the US and the Australian Incident Monitoring System). In healthcare there is a large variety of different reporting systems belonging to different agencies and institutions. Vincent (2006: 58) provides a list of examples of the different agencies including General Medical Council, Coroner, Health and Safety Executive, NHS Litigation Authority, Police, Nursing and Midwifery Council and so on. These all serve different purposes, for example, litigation and criminal investigation. Some of these systems (e.g., claims and litigation data) provide information for enhancing patient safety (e.g., the widely cited Harvard Medical Practice study reviewed closed claims data, Leape et al., 1991), but there is frequent duplication of function and confusion of purpose.

The Department of Health report *An Organisation with a Memory* (Department of Health, 2000) pointed out several shortcomings of reporting in the National Health Service (NHS). Subsequently, the National Patient Safety Agency was set up with a mission to implement the National Reporting and Learning System (NRLS) in order to bring about more coordination of information about patient safety issues and to produce wider dissemination of lessons from serious incidents. The primary aim of the system is described as to 'provide an independent system to record adverse events and near misses so that the NHS could minimise such incidents' (Carruthers and Philip, 2006: 12). Key objectives of NRLS are, therefore, to provide an overview of the *extent* and the *nature* of harm within the NHS and to develop solutions on a national scale.

As opposed to ASRS (confidential system), NRLS had been set up as an anonymous system to encourage reporting and to provide a more representative picture of the extent of harm across the NHS. In order to assess the nature of harm, NRLS requires information about the factors contributing to incidents. Since an anonymous system does not allow the analyst to follow up incident reports, NRLS includes a set of questions about contributory factors that are to be filled in directly by the reporter (see Table 17.1). The reporting process may include up to six different steps and some of the details which the reporter should fill in are related to the *where*, *what* and *how*, with a level of complexity that well reflects the healthcare domain (departments and specialities, phase of care, roles involved, etc.).

At the end of 2006, the Department of Health issued a report called '*Safety First*' to reflect on the past experiences (Carruthers and Philip, 2006). According to this report, the NRLS cannot be considered a success story. 'Despite the high volume of incident reports collected by the NPSA to date, there is little evidence that these have resulted in actionable learning for local NHS organisations. The NRLS is not yet delivering high-quality, routinely available information on patterns, trends and underlying causes of harm to patients' (p. 25). Such a negative verdict comes for an otherwise in many respects admirable approach that expanded on the FAA's experience to include state-of-the-art theories of organisational safety, such as those of James Reason (Reason, 1990, 1997).

Table 17.1 Table from the NRLS, with a list of contributing factors

ID06	What were the apparent contributing factors? (Tick any that apply)	
	□	*Communication factors* (includes verbal, written and non-verbal between individuals, teams and/or organisations)
	□	*Education and training factors* (e.g., availability of training)
	□	*Equipment and resources factors* (e.g., clear machine displays, poor working order, size, placement, ease of use)
	□	*Medication factors* (where one or more drugs directly contributed to the incident)
	□	*Organisation and strategic factors* (e.g., organisational structure, contractor/agency use, culture)
	□	*Patient factors* (e.g., clinical conditions, social/physical/psychological factors, relationships)
	□	*Task factors* (includes work guidelines/procedures/policies, availability of decision making aids)
	□	*Team and social factors* (includes role definitions, leadership, support, and cultural factors)
	□	*Work and environmental factors* (e.g., poor/excess administration, physical environment, work load and hours of work, time pressures)

Why is NRLS experiencing such problems despite the efforts that went into its design and implementation? To some extent, a brief comparison with major structural characteristics within which the successful ASRS operates, provides some insights into the problems that arise from design decisions that were taken for NRLS.

First, the distinction between adverse events, near-misses and events of less significance is more difficult than in aviation. The definition of adverse event usually adopted (harm incurred by a patient stemming from the process of care rather than from the illness itself) implies a full understanding of the clinical situation of a patient (see section 'The Pass Criterion').

As a result, it is not surprising that the main categories of incidents identified through incident reporting are concerned with adverse events such as patient falls and adverse drug events. A patient fall is clearly identifiable and often the reporter is in a good position to provide an account of the factors that played a role. For example, during the one-year period April 2006–March 2007 a total number of 727,236 incidents were reported to the NRLS. Of these, 265,343 incidents belonged to the category of patient accidents (patient falls, etc.). The next largest categories were treatment/procedure (64,227) and medication (62,660). The category of clinical assessment – a major activity within healthcare – contains only 35,316 reports. In view of the above deliberations this may not be surprising. This suggests that the categories that do get reported are those that are observable and identifiable, but this provides only very selective insights into the dynamics behind adverse events in healthcare: a large part of situations that pose risk are not reported because they cannot be identified as reportable incidents by the workers.

Second, NRLS adopted anonymous reporting to encourage a higher number of reports and moved the identification of contributory factors to a taxonomy within the reporting system to be filled in by the reporter rather than by the analyst. This was done in order to meet the dual aim of assessing the *extent* and the *nature* of harm within the NHS. For many events in healthcare, the relevant patient journey may span several shifts or even days and weeks and frequently the reporter is in no position to describe adequately the contributory factors without a thorough investigation, which is a clearly inappropriate task for the reporter. For example, an adverse drug reaction may be detected by a nurse or a doctor on a ward, but some of the main contributory factors may be distant in time and space, such as the possible failure to record a drug allergy on part of the patient's GP (family doctor). Such a constellation, where there are many different actors involved and the relevant activities unfold over a prolonged period of time and distributed in space, pose almost insurmountable problems to attempts to generate meaningful learning with a reasonable amount of effort from incident reports about the dynamics behind adverse events (this characteristic will be later referred to as 'Visibility').

Third, the specialisation of the healthcare domain makes it difficult to maintain a body of investigators to centrally analyse all of the reports. One of the lessons learned in the NPSA review of NRLS was that more clinical and front-line expertise was needed to ensure quick and accurate screening and acting upon reports received (see section headed 'Understand the characteristics of your community').

Fourth, compared to the aviation world, the healthcare world is a lot less uniform as well as less standardised. Even if we consider for the sake of simplicity only the world of secondary care, where most of the efforts in patient safety and incident reporting have been (albeit most of the patients' contact is actually within primary care), major differences within this domain become quickly evident. Secondary care presents with an extraordinary range of diverse activities, such as the mostly routine, but sometimes highly unpredictable and hazardous activities within surgery, the inherently unpredictable and constantly changing world of emergency care, or hospital medicine where diseases may be masked, difficult to diagnose, the treatment risky and complicated by multiple co-morbidities (Vincent, 2006).

Handle with Care: All Reporting System are Different

We have seen through the aviation and healthcare examples how incident reporting systems are better understood as deeply intertwined with their respective domains, and we analysed some structural domain characteristics that can contribute to their success or failure (i.e., whether incidents are observable and identifiable and whether the reporter or the analyst are in a position to identify and characterise the contributory factors). Can we identify more general properties that are present also in other domains of application?

In the following, we move from the two properties described in the previous section to offer a more elaborate reflection on five key dimensions that should be analysed when implementing a reporting system. We also describe how these dimensions should inform the decision on which reporting system to adopt. No clear cut answer on what is the best option can be given. Domain structural characteristics are a key dimension to analyse, but the final choice will always depend on what you are going to use your system for. Incident reporting systems serve many different purposes and an incident reporting system does not address by itself all purposes. A careful selection has to be undertaken.

The Pass Criterion

The first domain characteristic to be analysed is related to how easy events to be reported can be told apart from negligible ones on the basis of factual information, that is, with no (or minimal) subjective judgement. In some domains, it is possible to provide front-line people with clear guidance on what should be reported, while in other domains front-line people should exercise their best judgement to assess whether the event deserves reporting, or whether it has no significance. For instance, if we expect accidents and incidents to be reported, then we would need to analyse whether these events can be easily distinguished one from the other, and from other categories of events.

In case a clear cut distinction cannot be drawn, other concepts may be used. For instance, front-line operators may be requested to report risks (i.e., hazardous situations), which seems to be a viable solution especially in those cases where outcomes are hard to observe and assess, or if doubts exist that operators would report events with any consequence. In any case, the pass criterion remains valid also for risk reporting systems: guidance should be given on what to report and operators should undergo specific training to correctly identify events to be reported.

The key decision that should be informed by this domain characteristic concerns the risk–accident *continuum*, meaning whether one organisation should try to implement an incident reporting system or a risk reporting one. The first option is a viable one only for those domains where accidents/incidents can be distinguished from other events on the basis of factual information, that is, when the outcome of an event can be factually appreciated, with no need of subjective judgement. Risk reporting systems should instead be preferred in all those cases where events cannot be distinguished on the basis of the outcome, making it hard for front-line people to tell which type of event they just witnessed. In these cases, front-line staff should be more rightfully asked for their (subjective) perception of risk.

Degree of Standardisation

The degree of standardisation that is typical of one domain affects the type of information needed to reconstruct one event. In a standardised world like aviation, context can be often assumed to remain stable and can be left implicit, taken for granted. For a truthful reconstruction only the main events and actions are needed, while the scene on which the events unfolded can be assumed to be standard. For instance, in the ASRS system, operators report mainly on the event itself, but can disregard mentioning most of the contextual information. On the contrary, other domains may need a major emphasis on the context, as contextual features may be very relevant to determine the actual unfolding of the event. So this decision should be made considering the degree of standardisation of the reference domain.

To better elaborate on this concept, we may derive some notions from literary critique, more precisely from the work of Burke (1969). Burke defines an event with five elements: what was done (the act), who did it (the agent), when or where it was done (the scene), how he did it (means and tools), and why (purpose). According to Burke, these five elements are required to describe an event. These elements can be used to differentiate which information is typically requested to reporters in an incident reporting system or in a risk reporting one. Incident reporting may be said to focus on the description of the agent and the act (what has happened is the most important piece of information), while risk reporting may demand focusing on the scene or on the tool (the context is more important than the specific outcome).

The key decision that follows from the analysis of the standardisation degree concerns what information is required to reconstruct an event. In a structured domain, scene, tools and purpose can be considered as stable, thus taken for granted and left implicit in an event description. This would not be the case in less structured ones. In these latter cases, the description of the actor and of their actions should be complemented with a certain amount of information of why those actions were performed, with which tools and in which context. Information on actor and actions are not enough to reconstruct the event in a satisfactorily manner.

The standardisation degree is often linked with the pass criterion, in the sense that standardised domains often warrant a clear definition of what a significant event is, while non-standardised domains treat every event as a separate case. This correlation between the two different characteristics is often present, but it does not follow by any necessity. We may theoretically give the case of a non structured domain with a clear pass criterion, or vice versa. However, standardisation often affects not only the way operations are conducted, but also the expected outcomes, thus making it often proceed paired with the pass criterion.

Visibility

Once all of the above points have been scrutinised, we need to analyse operators' perspective in a realistic manner, in order to understand what operators are willing and able to report. Whatever our decisions have been on the other dimensions, the end question to address would be: Are the operators in a good position to observe the events I would like to collect? Are they willing to report?

A reporting system starts from the assumption that operators are an essential source of information. At this stage, we would need to challenge this assumption and delve further into it to understand in which respects and to what extent this is true. A similar recommendation comes from Eurocontrol, which advises to complement incident reporting with routine safety surveys to collect operators' feedback on their daily risk perception (Eurocontrol, 2000). Not all aspects of our system have the same degree of visibility, some may be easily perceived by front-line staff, whilst others may be hard to appreciate from their perspective. In other words, organisational processes may be shaped by visible factors, as well as by non-visible ones.

Visibility is affected by several dimensions, including how organisational processes are designed, the position and role of front-line staff, the duration and spatial span of processes, etc. It is also affected by the nature of the actual content of work. For instance, in the healthcare domain the content of work is the care of humans, and a human body has its specific dynamics, often not visible unless dedicated diagnostic activities are undertaken. As a result, not all the actions done on patients have an immediate, easy to appreciate effect, and some of the outcomes may be actually shaped by factors hard to single out. Moreover, different actors will most likely be able to perceive different aspects of the system, and will also possess different terminology and analytical skills to draft a report on that.

Considerations on visibility should inform the decision on what front-line people are asked to report. For each event to be reported, we should then analyse which aspects of the event are visible and which are not. For instance, the degree of visibility of aspects such as the outcome, causal factors, contributing factors, may be very different. This consideration should be complemented by the analysis on what front-line people are willing to report, which often depends on their safety culture.

Understand the Characteristics of your Community

The reference domain should be analysed also in terms of communities of practice and professional communities (Lave and Wenger, 1991). Micro-communities are likely to have different understandings (to a various extent) of the same situations, to appreciate different aspects, to use different tools and pursue different objectives (sometimes converging with other micro-communities, other times even contradictory). The more varied a domain as far as communities are concerned, the harder to establish a domain-wide (or nation-wide) reporting system. In other words, the immediate consequence of the heterogeneity of communities in one domain is on the scale of the reporting system. If the community is homogeneous, it is going to be easier to establish a large program, covering the whole domain or the whole country. On the contrary, very diverse communities may suggest establishment of more local systems.

Two other key decisions can be informed by the analysis of the community characteristics. First, analysts of incident reports should cover the whole spectrum of expertise (as in the ASRS case), to provide meaningful results and to competently analyse the information contained in the reports. If we get back to the discussion of the importance of context we just developed, analysts need to possess the background knowledge required to 'fill in the gaps' in the reports, to understand what is implicit and what was taken for granted by reporters. Different communities demand for a 'as varied as they are' body of expert analysts, which may be hard to put together and maintain. Second, the analysis of the communities in our domain can also inform how to provide feedback to reporters. For homogeneous communities, a non-targeted message may suffice, which would not be the case for more varied communities. In both cases, the feedback loop should remain as close as possible to operators to obtain effective results, to be as quick as possible in reporting back to them. But while a homogeneous community may leave extra space (the community has its own means of circulating the feedback and once the feedback is out it will spread quickly to everyone), diversity in the micro-cultures requires the loop to be as quick as possible and as targeted as possible.

Assess Safety Culture

The safety culture of different domains (and organisations) can be classified on the five levels described by Westrum (1993) – from pathological to generative. Each level distinguishes a different way in which domain members approach safety issues, from cultures that do not see any value in safety-related activities, to cultures that see safety as an integral part of everything is done. As far as reporting systems are concerned, safety culture will affect many dimensions. To mention just the main ones, we may list the protection offered to those who report, the amount of education about safety (awareness of safety issues), the ability to perceive causes of incidents (Dekker, 2007; Reason, 1997).

The key decisions to be made after having assessed the domain safety culture are those already listed for the other structural characteristics, including the accident-risk *continuum*, which information front-line people should report, which information they are able and willing to report, the type of feedback that can be offered, and the type of analysis that should be carried out on the reports. In addition to these dimensions, the safety culture level is a primary information to decide whether the reporting system should be anonymous or not, and the degree of confidentiality offered by it. Higher safety culture levels may warrant 'open systems,' with disclosure of names, while lower safety culture levels need to offer confidentiality, or even anonymity, as a condition to encourage reporting.

What Happens When Key Structural Properties are Missing?

We discussed how the success of incidents reporting can depend on some key structural properties of the target domain. The underpinning idea of incident reporting systems is that often the accidents and incidents precursors are similar. A lot of learning can be generated by focusing on events that can potentially cause harm (i.e., incidents), rather than exclusively focusing on actual harmful events. In this way, more data points are available from which more robust learning about the dynamics behind adverse events can be extracted. This is illustrated in Figure 17.1 below. In Figure 17.1, safety is represented as a control problem, i.e., accidents happen in the area where variation is out of control. The incident boundary represents the area where incidents are identified and reported. Each event provides a window through which the driving forces of harmful events and corresponding contributory factors (represented as arrows) can be identified, understood and subsequently generalised across the range of incidents.

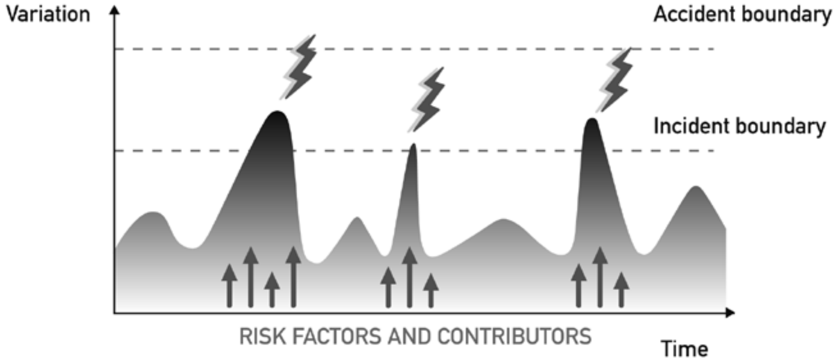

Figure 17.1 Safety represented as a control problem

For domains and organisations where the above structural characteristics for successful incident reporting systems do not hold, a promising approach may be to focus directly on the driving forces behind adverse events.

Proactive Risk Monitoring

This approach is inspired in part by Reason's Tripod methodology (Reason, 1997) developed for the oil and gas industry. Tripod suggests the monitoring of basic risk factors at regular intervals either through audit or through feedback from staff. In this way, a risk profile can be built up over time and basic risk factors most in need of addressing can be focused on (in Figure 17.2 the upward arrows represent forces that drive variation, the downward arrows represent forces that lead to a control of variation). In this way, we are reliant neither on incidents as triggers (which may not be easily observable) nor on the ability of a single reporter to provide a full account of a complex system dynamics.

A main difference to incident reporting is the fact that risks (or rather: factors contributing to risk) are monitored themselves. Reason identifies as basic risk factors organisational processes giving rise to latent conditions, such as processes for procurement of equipment, maintenance management, processes for defining communication interfaces etc. This is in line with Reason's model of organisational accidents that suggests that accidents are the result of multiple active and latent failures, where only the latter are sufficiently predictable and controllable. This approach to risk monitoring focuses in particular on the forces that drive variation out of control (upwards arrows in Figure 17.2). Alternative models, such as the Functional Resonance Analysis Method (FRAM) (Hollnagel, 2004) may give rise to a different emphasis. The concept of resilience as 'the capabilities on all levels of a system to respond to regular and irregular threats in a robust yet flexible manner, and to anticipate the consequences of disruptions' (Hollnagel et al., 2006) highlights how reporting system may be better aimed at detecting *near-resonance* situations, that is those situations where the system faces 'disruptions and variations that fall outside of the base mechanisms/model for being adaptive as defined in that system' (Woods, 2006a: 21).

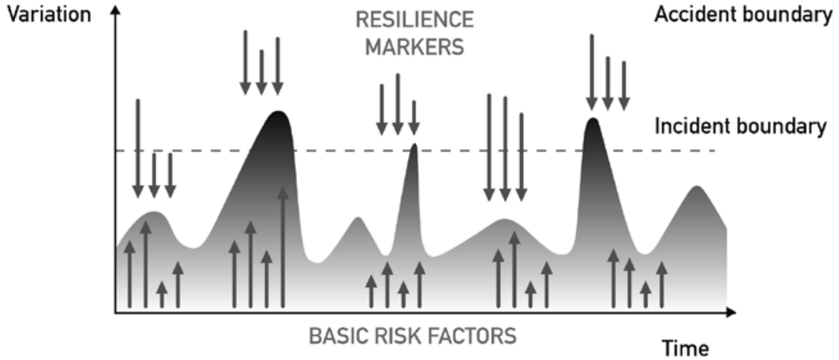

Figure 17.2 Risk monitoring elicits regular feedback from staff about a number of key contributory processes and factors

In either of these approaches, there is a shift from the concept of incident reporting to the identification, reporting and analysis of situations that fall outside the 'design envelope,' in order to better understand how 'a system is competent at designed-for-uncertainties' (Woods, 2006a). Instead of focusing on organisational processes that give rise to latent conditions, the aim of a risk monitoring system becomes the monitoring of variations within activities and processes or more generically the extent to which an organisation and staff are capable of anticipating, recognising and adapting to variations and disturbances. This approach emphasises also the positive side of performance by taking into account the forces that enhance control of variation (downwards arrows in Figure 17.2). Both, the establishment of tripod-like risk monitoring and the development of meaningful markers of resilience within healthcare are currently the object of ongoing research.

Conclusion

The aim of this chapter has been to show how reporting systems should be considered as tools embedded in socio-technical systems. A reporting scheme has no unique objective in it, but may serve different purposes. A task for any reporting system is then to clearly target some objectives and to design a corresponding process of data collection, analysis, feedback and action. Our aim is to provide an initial answer to the following research question: which are the structural domain and organisational properties we need to look for when trying to implement reactive and proactive risk monitoring systems? Which objectives can these systems address? By analysing the use of incident reporting systems in civil aviation and in healthcare, we reflect on their role within the wider safety management system. Not all organisations are equal, so each may require its own reporting system. To implement a reporting system, it is thus necessary to clearly target some objectives and to design the corresponding process of data collection, analysis, feedback and action.

This chapter started from a literature review to highlight some structural characteristics of a domain that should be considered when designing an incident reporting system. If we were to summarise the five structural characteristics we have discussed above and to find a common explanation for them, the best way would probably be to reason in terms of domain culture. From the above discussion we see that the aviation culture presents a good degree of homogeneity, which ensures stable definition of operations, of anomalies and of expertise. Even if micro-cultures are present (e.g., pilot community, air traffic controllers, cabin crew, etc.) these are well recognised and their voice is represented in the ASRS panel of experts, so that the community can speak with one 'non-controversial' voice. Billings (a founding father of the ASRS system) clearly states that one key requirement for a successful incident reporting system is 'a demonstrated, tangible, *widely agreed upon* need for more and better information' (Cook et al., 1998: 52, emphasis added). Billings also states that *consensus* is not enough and that *understanding* of what the ASRS is doing is a necessary point, among all the stakeholders (p. 55). Both consensus and understanding can be considered as indicators of a shared culture in the aviation community. The healthcare domain does not exhibit a comparable degree of sharing and definitively presents a more varied array of different professional communities.

In the last section, we have linked the five structural characteristics to key decisions to be made when establishing a reporting scheme. We have also described how reporting systems can be shaped in different forms, once a good awareness of their purposes has been achieved. This variability has been placed on a three-dimensional continuum:

- from risks to accidents
- from the scene to the event
- from small scale to nation-wide scale.

Being linked with structural characteristics, these dimensions are to some extent domain independent, and can be used to compare systems across various domains.

Chapter 18
Is the Aviation Industry Ready for Resilience? Mapping Human Factors Assumptions across the Aviation Sector

Kyla Zimmermann, Jean Pariès, René Amalberti and Daniel H. Hummerdal

This research maps out differences in safety perspectives, the Traditional and Resilient, respectively, across the civil aviation industry in Europe and the Americas. We surveyed 705 aviation professionals to determine whether they agreed or disagreed with the human factors and safety assumptions currently dominant in the aviation industry's tools and methods. The results show variations in perspectives according to the national culture and occupation of the respondent. We also discovered that non-experts in human factors (HF), ergonomics, or safety may indiscriminately and unconsciously use HF ideas from a variety of (sometimes conflicting) safety models or paradigms. The results of this study can help Resilience Engineering researchers and safety managers to better understand the point-of-view of practitioners who use their tools and models.

Paradigms in Safety and Human Factors

Resilience Engineering is the culmination of over 25 years of accumulating evidence supporting a different way of thinking about organisational behaviour and the performance of complex, high risk, socio-technical systems (Amalberti, 2001; Hollnagel, 1983; Hollnagel and Woods, 1983; Pariès, 1996, 1999; Perrow, 1984; Turner, 1978). This way of thinking is not only different; it directly conflicts with some of the fundamental assumptions which define human factors or ergonomics theories applied in industry today. Thus Resilience is the antithesis of the traditional and still prevailing, human factors and safety paradigm, referred to by Hollnagel as the 'Traditional Safety Perspective' (see Table 18.1).

Models and methods based on these assumptions may not meet the needs of ultra-safe, complex, modern industries such as aviation and may prevent further progress (Amalberti, 2001, 2006; Dekker, 2005; Hollnagel, 2004; Le Coze, 2005; Leveson, 2002; Marais et al., 2004; Pariès, 1999; Rasmussen, 1997; Woods, 2004). There is a lot of academic momentum building in favour of change yet 25 years on, the traditional ideas seem to remain entrenched in the perspectives and approaches of industry practitioners.

This is not a simple transition to make, particularly at a large scale (Steele and Pariès, 2007). Amalberti explains that a mature system such as commercial aviation may no longer have the flexibility for dramatic or profound change (2006).

Table 18.1 Contrasting system perspectives (adapted from Hollnagel, 2008)

Technological Optimism ('Traditional Safety Perspective')	**Technological Realism ('Resilience Engineering Perspective')**
Humans are a liability. Variability is a threat to safety and efficiency. Design should constrain variability Things go right because: • systems are well designed and scrupulously maintained; • procedures are complete and correct; • people behave as they are expected to – as they are taught; • designers can foresee and anticipate every contingency.	Humans are an asset. Humans are necessary for technical systems to function properly. Things go right because people: • learn to overcome design flaws and functional glitches; • adapt their performance to meet demands; • interpret and apply procedures to match conditions; • can detect and correct when things go wrong.

Figure 18.1 is an overview of the evolution of the existing human factors and safety science paradigm as well as the aspirations of Resilience Engineering, adapted partly from the Resilience symposia and literature (Hollnagel et al., 2006). Normally this table is presented as three distinct columns (Hollnagel, 2004) but in fact the underlying scientific paradigm of the epidemiological model (the Swiss cheese model) is just an extension of the sequential model (dominoes). Resilience is more of a revolution than evolution, breaking with the past at an epistemological level.

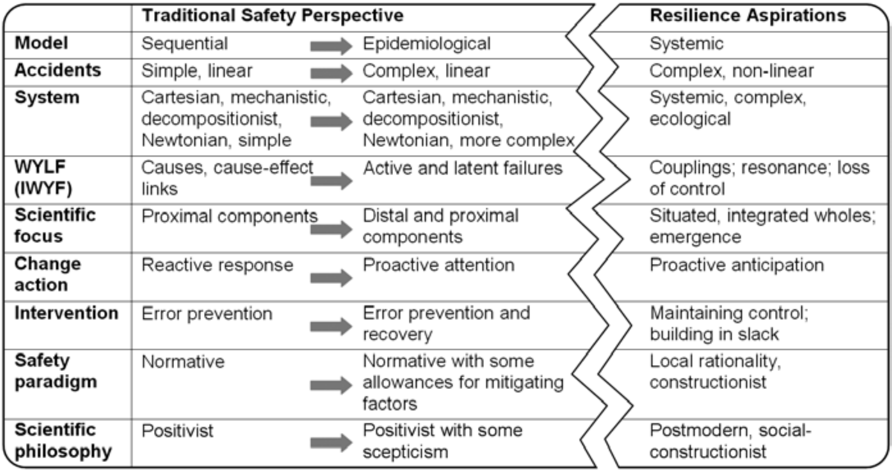

	Traditional Safety Perspective			Resilience Aspirations
Model	Sequential	→	Epidemiological	Systemic
Accidents	Simple, linear	→	Complex, linear	Complex, non-linear
System	Cartesian, mechanistic, decompositionist, Newtonian, simple	→	Cartesian, mechanistic, decompositionist, Newtonian, more complex	Systemic, complex, ecological
WYLF (IWYF)	Causes, cause-effect links	→	Active and latent failures	Couplings; resonance; loss of control
Scientific focus	Proximal components	→	Distal and proximal components	Situated, integrated wholes; emergence
Change action	Reactive response	→	Proactive attention	Proactive anticipation
Intervention	Error prevention	→	Error prevention and recovery	Maintaining control; building in slack
Safety paradigm	Normative	→	Normative with some allowances for mitigating factors	Local rationality, constructionist
Scientific philosophy	Positivist	→	Positivist with some scepticism	Postmodern, social-constructionist

Figure 18.1 The traditional safety and resilience engineering models as applied to aviation systems and safety

Diversity in Aviation

Commercial aviation is highly standardised and regulated at an international scale. However, as these different safety paradigms demonstrate, there is still room for interpretation and variation of how people perform, understand and manage the work. We undertook to characterise the variation in perspectives about human factors and safety across the aviation community, specifically the agreement with the Traditional Safety or Resilience Engineering perspectives.

The Safety Assumptions and Resilient Attitudes (SARA) Survey

The objective of the research was to map the differences in safety perspectives held by aviation professionals across the industry in Europe and the Americas. We were primarily interested in the variation across the aviation professions, but we looked at 20 factors, including national culture, HF background, and work experience.

The Business of (not) Measuring Resilience

We set out to ask questions to gauge who in the industry used Resilience Engineering as a theoretical framework. We were immediately confronted with practical problems: First, resilience is still more easily defined (even by the Resilience Engineering community itself) in terms of what it is not than what it is. Second, many Resilience Engineering ideas appear complex and do not lend themselves easily to survey methods. Finally, presenting Resilience ideas, like those in the right-hand column of Figure 18.1, in a survey would not be meaningful – they seem like common sense (even though they are far from the industry norm). So our solution was instead to gauge the level of agreement with the contested Traditional assumptions still so prevalent in the news media, accident reports, and safety media. These assumptions, like those in the left-hand column of Figure 18.1, help define Resilience indirectly. Thus our aim was to infer Resilience Engineering attitudes by the rejection/acceptance of the Traditional Safety perspective.

Mapping the variations in perspectives across the industry would allow safety and HF professionals interested in Resilience to better understand the practitioners – the target audience of their work. Assumptions and beliefs form the safety paradigm, which in turn lays the foundation for safety action through models, tools, etc. (Lundberg et al., 2009). There is more than one way of seeing 'the facts' and this is influenced by one's paradigm, whether or not one is aware of it (Gergen, 1999; Simpson, 1996). Hence making this paradigm explicit is useful, especially as Resilience researchers are pushing for changes. Mapping Safety Assumptions and Resilient Attitudes (SARA) results across the industry would tell us whether aviation is ready to make the paradigm shift to Resilience.

Developing the SARA Survey

In the first phase of the research, we gathered examples of frequently contested assumptions from the literature, such as those in Figure 18.1 and in Dekker's *Ten Questions on Human Error* (2005), which define the existing Traditional aviation safety paradigm (similar to Dekker's 'old view').

We developed a series of questionnaire items based on these assumptions, and surveyed aviation professionals in different domains and countries to find out whether or not they agreed with the 'Traditional Safety Perspective' on a scale of 1 to 5. We also conducted an initial round of confidential interviews with aviation professionals representing different jobs and geographic regions to assist and inform the survey writing process.

The Survey Respondents and Interview Participants

Using the 'snowball' distribution technique, the anonymous on-line survey was disseminated in three languages to a convenience sample based on the researchers' professional networks and the partners of the European aviation research project HILAS (Human Integration into the Lifecycle of Aviation Systems). Tables 18.2 and 18.3 show the distribution of the 705 survey respondents according to job and geographic region. We grouped countries into geographic regions based on culturally similar clusters identified in the GLOBE Study (House et al., 2004).

Table 18.2 **Distribution of the 705 survey responses according to geographic region**

Geographic region	%
Anglo (English-speaking nations): UK, US, Canada, Australia, New Zealand, and South Africa	29
France	25
Northern Europe (EU-N): Scandinavia, The Netherlands, and Flemish-speaking Belgium	11
Southern Europe (EU-S): Spain, Portugal, and Italy	10
Latin America (Latin Am): Mexico, the Caribbean, South and Central America	9
Eastern Europe (EU-E): the new EU members and CIS states (except Russia), Turkey, Greece, and Malta. France also includes responses submitted in French from Switzerland, Luxembourg, and Belgium	8
Other	8

Based on the results of the statistical analysis of the questionnaire, we carried out a second round of confidential interviews with participants from specific regions and occupations.

Analysis of the SARA Survey Results

The survey data were analysed using qualitative statistics (multiple correspondence analysis or MCA) to allow the significant factors to emerge on their own without the *a priori* influence of the chosen analysis criteria. We will present the two factors which predominantly explained the variation in questionnaire response using simple descriptive statistics (e.g., averages) since they reveal the same result as the MCA but are more familiar and simpler to interpret. In order to analyse the data quantitatively, a ‘score’ was calculated for each respondent by averaging their ratings on the five-point scale for the 33 questionnaire items. The score was used as a measure of the respondent’s agreement with the Traditional Perspective and is the basic unit of comparison for the analysis described in this chapter. The differences discussed here are all statistically significant ($p< 0.05$ or better).

Table 18.3 Distribution of the 705 survey responses according to job

Job type	%
Other job	27
Pilot	26
Air traffic controller (ATCo)	16
Design or other types of engineer (Engineer)	15
Aircraft mechanic, technician or maintenance engineer (AME)	8
HF or ergonomics expert or safety manager (HF/Ergo/Safety)	8

While necessary for the analysis, the idea of a ‘score’ is inherently contradictory to our research philosophy: it is not our intention to judge people as wrong or misguided based on whether they agree or disagree with us. It is our hope that readers of this book share this philosophy and will not misuse the data. The real potential in these findings is to highlight the variations in opinions across the industry so that researchers and safety managers can better design and target safety interventions for these diverse populations.

Differences between National and Occupational Cultures

The largest variation in the survey results was according to the country of residence (Figure 18.2), followed by differences among the professions (Figure 18.3). In both figures, the centre of the bar represents the mean score while the total length of the bar represents 68.2 per cent of the scores (two standard deviations) for normally distributed data. The respondents who most strongly disagreed with the traditional perspective were those living in Northern Europe or working as HF researchers/specialists. The respondents who were most in agreement with the traditional perspective were those living in Latin America or Southern Europe (Italy, Spain, and Portugal), or aircraft mechanics (AMEs). There was no interaction of these two factors. Overall the averages of all scores appear close to neutral on the scale, however the differences shown on the graphs are statistically significant and show an interesting trend: disagreement with the assumptions is roughly correlated with latitude (Hofstede, 2001), with the usual exceptions of New Zealand, Australia, and South Africa (part of the Anglo sample).

The follow-up interviews were done with pilots, ATCos, and AMEs from countries in the Northern and Southern European regions in order to represent the largest gaps among practitioners. Although the professions exhibiting the largest differences were AMEs and HF/Ergo/Safety, we did not interview the latter because we had never actually intended to survey a significant number of our colleagues; their views are already a matter of public record. However, this proved fortuitous as a measure of the survey's validity (many of the HF/Ergo/Safety respondents did not answer anonymously, and we know them to be opponents to the Traditional Perspective).

During the interviews in Southern Europe, we had to re-examine the logic of clustering Spanish and Italian responses after several Spanish participants claimed their perspective on aviation safety to be more similar to that of France than Italy. Closer examination of the data confirmed this to be true for the Spanish (n=19) SARA scores. When analysed alone, the responses from Italy (n = 48) were further to the Traditional end of the scale, more similar to the Latin American sample.

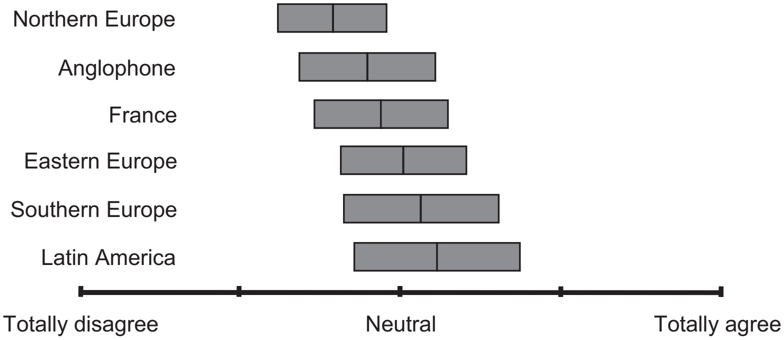

Figure 18.2 Average survey responses according to geographic region

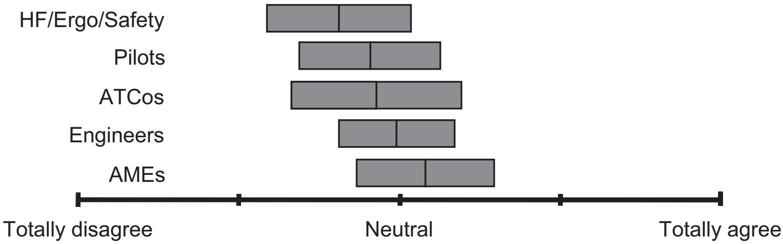

Figure 18.3 Average survey responses according to job

Discussion of Differences between National/Societal Cultures

Given the difficult business of quantitatively measuring ‘safety’ and defining the limits of societal cultures (a more correct term, as culture does not necessarily correspond to national borders) it is not surprising that there is very little available material for comparing safety outcomes across nations and making conclusive links to culture; there is ‘no smoking gun’, as Hutchins et al. put it (2002: 6). What does exist is controversial and incomplete, showing only crude comparisons between continents (Civil Aviation Authority, 1998) and these data is highly contested due to confounds in the metrics. It is not informative for the purposes of our study, unless we compare only the continents of North and Latin America, Western and Eastern Europe. In that case we can see a correlation between higher SARA scores and quantity of aircraft hull losses.

There is non-aviation safety data available (e.g., road safety or occupational accidents) tending to support the stereotype that the North, West and Germanic European countries (considered more rule-based, law-abiding, orderly and stoic) have better safety outcomes than their Mediterranean or Latin European neighbours (considered more *laissez-faire* about rules and more passionate or emotional about life) or the new EU member states or CIS states (Zimmermann, 2009). Europe is the most culturally diverse region of the world. Comparing Scandinavia and Italy reveals a very different profile along Hofstede’s cultural dimensions. Yet politically things have changed so dramatically throughout the 20th century that cultural assessment is a moving target (e.g., the Spanish SARA results mentioned previously). It is difficult to rely even on recent, large-scale studies such as those by House et al. (2004) or Hofstede (2001), as they may already be out of date in some respects.

The Limitations of Understanding Culture

Comparisons using surveys and interviews are problematic in themselves since the semantics of terms is subjective and might have cultural or individually determined connotations. For example, we overheard a German colleague complaining that trains in Switzerland are 'always late'. This is a salient example of the subjectivity of even concrete concepts like 'late' and 'always' – so how can we have meaningful discussions about 'safety' or 'risk', etc? Hutchins et al. explain that 'at first glance, the effects of national culture appear pervasive and obvious, but when one seeks a theory … or when one looks for direct evidence of the effects of culture …, culture seems to vanish' (2002: 6). This is one of the dilemmas of cultural research, and makes it difficult to explain the real meaning behind the different SARA scores. However, since our objective was mainly to map the differences for practical purposes, understanding the different interpretations of the questionnaire items is not essential (but it does provide interesting opportunities for future research).

Behaviour and Attitudes Understood Within their Context

During the follow-up interviews participants were quite moderate in their praise or criticism of other continents or regions within Europe in terms of safety or rule-following behaviour. However, there were some evocations of the North–South stereotype mentioned above. Interestingly most participants described their own culture as the ideal balance between rule following and creativity and considered those to the North as slightly too rigid, and those to the South as slightly too unpredictable (Zimmermann, 2009). There may be some truth to this; it could be a reflection of the clashes occurring when people encounter different operational cultures that do not match their expectations. The behaviours may be different but still appropriate for the local context.

While the orderly, rule-following cultures may seem logically safer, that depends on the scenarios, time-scale and measures chosen to assess 'safety'. The Tenerife disaster is an example of a paradoxical negative side-effect of safety rules and how the pressure of a strict rule-following culture can pose a different sort of risk.

Resilience and/or Safety?

One participant from Southern Europe explained his point-of-view on the difference approaches as they relate to aircraft maintenance:

> In Northern Europe they won't do the work until they have all the equipment and tools and parts and things. That's fine for them because they can get everything they need. We [in Southern Europe] don't have everything we need – sometimes we don't have anything – so we can't check off the boxes … the way we're supposed to, but we get the job done anyway, … use creativity … In the North they would just stop working. We can't work that way or we could never get anything done. (Zimmermann, 2009)

This raises the question of Resilience. One might argue that the individuals who are used to coping in 'adverse conditions' like those described above may actually be more resilient; the challenging working conditions offer more opportunities to practice their skills and develop problem solving. The work environment may also be more loosely coupled by necessity, making it more flexible in the face of crises. However, if these resilient properties exist at the micro-level only because the macro-level system is stretching things too thin (e.g., not adequately supporting the work or not providing the needed tools and infrastructure) then the system as a whole would be unprepared to deal with problems for a multitude of other reasons.

An Italian interviewee recounted an ironic anecdote:

> Of course Northern Italians feel that driving in Southern Italy is more dangerous, but the fact that no one stops at a red light means that everyone is paying attention. In the North, people assume if the light is green, they can go through without looking, so actually it is more dangerous to drive in the North. (Zimmermann, 2009)

Because of constant exposure to a dangerous environment this Southern Italian driver uses Simpson's *Cautious Cognitive Framework*: everything is considered a hazard unless it is explicitly indicated otherwise (1996). The irony is evident here; that the driver himself presumably also ignores traffic signals – so the net result may not be objectively safer. We know from HRO theory that it is a challenge to maintain that Cautious framework (keeping the possibility of risk alive) in a system with few accidents or incidents.

These two examples from the interviews about national/societal culture illustrate a paradox of the relationship between Resilience and safety: An unsafe system may be more flexible, more cautious, and may inadvertently foster Resilience at the micro-level. Similarly, a stable, safe system would have difficulty maintaining it. As Rasmussen (1997) has shown, when things go well the natural tendency is to increase production levels. This could increase the inherent risk (e.g., more planes or passengers), reduce the flexibility, tighten the couplings, etc. Amalberti (2006) describes the different types of actors (e.g., pioneers, craftsmen, and equivalent actors) in systems with different levels of safety. As aviation keeps evolving towards higher levels of standardisation, automation, procedures, and stability we must recognise that this comes at the expense of Resilience (Holling, 1973).

Discussion of the Differences between Occupational Cultures

Professions often create and sustain their own specific cultures – the training, the environment, the expectations of others, as well as the nature of the work itself influences and is influenced by the traits, skills, attitudes and behaviours of its members. There are common stereotypes about professions (although these are to some extent specific to a the local culture) and at a dinner party one would probably respond differently upon meeting a primary school teacher or politician, a bank teller or a Rock 'n Roll musician, a cashier or a university professor. Within the aviation community, like any other, there are stereotypes highlighting the differences between pilots, cabin crew, managers, inspectors, ATCos, etc.

Although Hofstede's (2001) seminal work is most known for his characterisation of national culture, he points out that the differences between professions in his results are also very as significant. For example, his sample of senior managers across the world had more in common with each other than senior managers and secretaries from the same country. This evidence supports the concept of occupational culture, demonstrating its significant influence on values and way of thinking.

Helmreich and Merritt (1998) found that pilots do score differently than their country averages on some Hofstede's cultural dimensions, and Hutchins et al. (2002) use pilots as an example of professional culture. Lumpé (2008) did a study on airlines in Europe and explored the idea even more deeply, explaining that the various professional cultures within an airline were so different as to merit unique management and leadership strategies.

Considering the results of our study combined with this evidence from previously published work, HF researchers and safety managers should be able to appreciate that their way of thinking about safety and HF topics differs from that of practitioners, and in particular AMEs. Additionally, Safety and HF interventions may need to be tailored to take into account not only the target national/societal culture but also the profession, since a standard 'one size fits all' approach will not be as effective.

Cultural Bias in Culture Research

Human Factors, like many aspects of aviation, is a cultural artefact based on the Western way of thinking. It is possible that the definition of the Resilience paradigm itself is culturally bound and may be incompatible with certain cultures. In this study Northern Europeans and HF specialists most strongly rejected the Traditional Safety Perspective. This raises the question of whether our tool is biased because our background and approach are rooted in these two areas to a large extent. We offer this deliberate misquotation of Dekker and Woods (2002) as food for thought: If Resilience gets to pick the battlefield, Resilience will win.

Ambiguity and Contradictions

Three other results are worth mentioning here and analysed together they may offer more insight to HF experts and researchers on how to improve resilience in the aviation industry.

The first result was that the means of all the groups within the sample were near the centre of the scale, expressing a neutral opinion towards the assumptions in the survey. Second, there was a large variation within the results of each individual respondent (but our analysis showed the results were not random).

Third, during the interviews we had hoped to reveal the safety perspective of the participant, but instead we observed that in terms of a safety philosophy or paradigm, participants radically contradicted themselves throughout the interview. For example, one respondent claimed that 'following the rules assures total safety', then gave an example of a problematic company rule and explained that 'following the rules does not mean you'll be safe'. It was not unusual for a single participant to express attitudes and beliefs representing the range of paradigms derived from the HF literature.

During the course of our research we also observed an airline revising their internal incident investigation process. Their new manual proscribed the application of several different, independent investigation methods in spite of some inconsistencies between them. All of the methods were well known industry tools. Although this approach seems illogical to us, from the company's point of view each of the methods must be 'right' because it was published somewhere by an expert, thus applying multiple methods will give results that are 'even more right'. They never questioned the underlying model or why traditional incident investigation techniques were falling short of the airline's needs.

Discussion: Integrating and Interpreting Ambiguity and Contradictions

Considering these findings together indicates to us that practitioners (i.e., non-experts in human factors or ergonomics, etc.) may not have a coherent, consistent, complete framework guiding how they view and understand safety. They may call up individual ideas from different paradigms or frameworks depending on the situation or the cognitive availability of the idea. There are many possible explanations for this, among them that practitioners may not have or need a coherent framework and may not even be aware when they express contradictory ideas. Or they may realise that there are frameworks but may apply different ones to different situations. As Hollnagel (2004) and Amalberti (2006) have suggested, different accident models or approaches may be appropriate in different contexts.

The Limitations of Attitude Measurement

These results also raise the problem of whether it is even possible to 'measure' beliefs or attitudes in this way. Attitudes are assumed consistent beliefs held by an individual. From the literature we learn that people's attitudes in response to a certain question may change quite a lot depending on the context, how a question is phrased, or even the order of questions in a survey. Repeated measures of attitudes shows low stability in how people respond in surveys (Bertrand and Mullainathan, 2001; Zaller and Feldman, 1992). Burr (2003) argues that attitudes are neither stable nor coherent. It appears rather as if ideas (and attitudes) are in the service of action (Schwarz, 2007). Talk and behaviour are some of the tools people use to bring about different effects and points in their social encounters. That means that people may change their responses depending on, for example, how they want to portray themselves in the survey or interview (known as 'social desirability bias'), or depending on what other ideas or desired actions were triggered by the question (e.g., to distance themselves from/identify with a colleague who 'made a mistake').

For researchers the paradigm or scientific framework is explicit as part of our professional identity. We normally need to have a well defined, theoretically consistent, stable perspective in order to do our work. Whereas researchers may see competing and/or contradictory ideas, practitioners may not. From the field of conversation analysis it has been shown that an individual may use completely contradictory ideas depending on what point he or she is trying to make (Potter and Wetherell, 1987). For example, in our interview data, it was common for a participant to express beliefs that accidents may arise from normal daily operations and environmental constraints, as well as claiming that accidents stem from some remarkable human error or failure, deserving of inquiry and sanctions.

Different discourses about what is safe and what are sources of risks are available for practitioners to position themselves along different dimensions or define their identities. They may draw from any convenient paradigm or popular folk wisdom that seems applicable to justify an action or explain a phenomenon. Hence, variances in the local context of culture, groups and individuals may introduce conflicting responses. As people become aware that there are competing models (as HF experts and researchers supposedly are) they may be less likely to express such contradictory ideas.

Two Explanations, One Conclusion

The neutral overall response combined with a large internal variation in scores could have another obvious explanation: the survey may not do its job, it might lack validity and not measure what it claims to. However, even if this is the case the implications of the result are the same since the questionnaire items represent common assumptions prevalent in accident reports and safety literature taken out of context (such as 'human error is the largest threat to flight safety'). Thus if the survey is measuring what we intended, it indicates that practitioners do not have a strong allegiance to the Traditional Perspective. If the survey's validity is in doubt, this demonstrates how the common HF assumptions or 'folk wisdom' used indiscriminately in aviation safety are ambiguous and even unconvincing when presented out of context.

In either case, the bottom line is that the predominant HF paradigm does not act as a successful framework for practitioners to understand non-technical issues in a consistent manner. Although not loyal to any single perspective, at least practitioners appear open to plausible-sounding argument. If Resilience hopes to do better we have our work cut out for us to create models and tools which are clear enough for users to recall and apply in their daily work.

Is Resilience Ready for the Aviation Industry?

We started out by asking whether the aviation industry is ready for Resilience. We were curious whether they are prepared and motivated to accept the Resilience perspective mainstream. During our inquiry of the safety perspectives of a sample of aviation professionals from around the Western world we discovered that the answer lies partly in the question: Is Resilience ready for the aviation industry? The insight gained from this study can help Resilience researchers and safety managers bring about the much-needed paradigm shift in aviation.

First, we saw that HF researchers and safety leaders (who develop the systems, model, rules and safety management tools) need to consider their target user population, since our survey results revealed variations in perspectives according to national culture and occupation. Second, the aviation practitioners in our study were not loyal to the Traditional Safety Perspective – in fact they may not think about HF and safety using any coherent perspective the way HF specialists do. This implies both a problem and an opportunity for Resilience advocates: It may be easy to convince people that Resilience makes sense, but they may not necessarily apply it consistently or exclusively.

The variations in perspectives identified in this study also help to dispel the myth that aviation is a purely technical domain in which standardisation has eliminated all variations in how people do the work. The human contribution is critical since there is always room for interpretation, no matter how standardised and regulated it becomes. Flying, controlling, and maintaining aircraft involves more than just checklists, radio frequencies, and torque settings. It took tragedy for the world to recognise how national culture influences communication on the flight deck and with ATC. Likewise, cultural differences (of any type) cannot be ignored, as they exist throughout the commercial aviation industry.

Acknowledgements

This work was supported by Dédale, the Foundation for an Industrial Safety Culture (FonCSI), and HILAS 'Human Integration into the Lifecycle of Aviation Systems' (a European Commission 6th Framework project).

Epilogue: RAG – The Resilience Analysis Grid

Erik Hollnagel

Resilience is defined as the intrinsic ability of a system to adjust its functioning prior to, during, or following changes and disturbances, so that it can sustain required operations under both expected and unexpected conditions. Since resilience refers to something a system does (a capability or a process) rather than to something the system has (a product), it cannot be measured by counting specific outcomes such as accidents or incidents. This chapter describes an approach to measure the resilience of a system that focuses on the four main abilities that together constitute resilience: the ability to respond, the ability to monitor, the ability to anticipate and the ability to learn. These abilities can be assessed by means of a number of questions and the answers can be represented by an easily comprehensible graphical form. This can be used to compare consecutive measurements, and thereby as a way to support the management of a system's resilience.

Introduction

Resilience Engineering is concerned not only with *what* makes systems resilient and how to make them resilient (i.e., to *engineer* resilience), but also with how to *maintain* or *manage* the resilience of a system. Since resilience refers to a quality rather than a quantity, to something that a system *does* rather than to something that a system *has*, managing resilience can be seen as a kind of process control. In order to manage or control a process, whether it is the resilience of a system or the steering of a vessel from point A to point B in an archipelago, three things are necessary. It is first of all necessary to know the current status or present position, that is, where one actually is at the moment. Second, it is necessary to know what the goal is, to have a clear idea about what the *future* status or position should be and therefore to know in which direction to move (using a Euclidean metaphor). And finally, it is necessary to know how a change can be made, specifically a change in direction, in magnitude, in speed, etc. It is, in other words, necessary to know the *means* by which a specific change can be brought about. While all three are essential for effectively managing a system's resilience, this chapter will focus on the first.

Resilience versus Safety

A system is usually considered safe if the number of adverse outcomes can be kept acceptably low. This can mean the accidents and incidents that may happen, but can also include adverse outcomes of other types such as work-time injury, work-related illnesses, etc. The advantage of defining safety in this manner is that the level of safety can be measured by counting various types of outcomes. The common understanding is, of course, that a higher level of safety corresponds to fewer adverse outcomes – and *vice versa*. One example of that is the International Civil Aviation Organisation's definition of safety as 'the state in which the risk of harm to persons or of property damage is reduced to, and maintained at or below, an acceptable level through a continuing process of hazard identification and risk management' (ICAO, 2006: 1). In a similar vein, the Patient Safety Indicator guide published by the US AHQR (Agency for Healthcare Research and Quality) defines safety as 'freedom from accidental injury', or 'avoiding injuries or harm to patients from care that is intended to help them'. In other words, safety is defined by the absence of adverse outcomes.

There is, however, more to safety than just reducing the number of adverse events. Resilience Engineering argues that it is necessary to focus on what can go right as well as at what can go wrong. From a Resilience Engineering perspective, failures arise from the adjustments needed to cope with the underspecification of the real world rather than from a breakdown or malfunctioning of normal system functions. Being a practical discipline, Resilience Engineering therefore, looks for ways to enhance the ability of systems to continue to function in as many different situations as possible, and safety is consequently defined as *the ability to succeed under varying conditions*. This definition includes the traditional meaning of safety, both because failures will adversely affect the ability to succeed and because an increase in the number of things that go right means a decrease in the number of things that go wrong. But the definition of safety also focuses on the system's ability to function under varying conditions, with consequences for how resilience is measured and for how it is managed.

Reactive and Proactive Adjustments

The key feature of a resilient system is its ability to adjust its performance. Adjustments to how things are done can, in principle, be *reactive* and take place *after* something has happened, be *concurrent* and take place *while* something happens, or be *proactive* and take place *before* something happens.

- *Reactive adjustments* are by far the most common. They happen in the aftermath of an event, for instance following the recommendations issued after an accident investigation, or as the 'lessons learned' from a major change or disruption. Responding when something has happened can, however, not guarantee a system's safety and survivability, even if the response is fast. This is because a system can only be prepared to respond immediately to a limited set of events. Since it will take longer to respond to all other events, the response is less likely to be effective.
- *Concurrent adjustments* are basically fast reactive adjustments that take place while a situation is still developing. For instance, if there is a major accident in a community, such as a large fire or an explosion, local hospitals will change their state of functioning and prepare for the rush of people that may have been hurt (Cook and Nemeth,

2006). Concurrent adjustments are the basis for continuous regulation, as described by the common feedback control loop (cf., Chapter 5).

- *Proactive adjustments* mean that the system can change from a state of normal operation to a state of heightened readiness, and possibly also act, *before* something happens. A state of readiness means that resources are allocated to match the needs of the expected event, that special functions are activated, that defences are increased, etc. An everyday example from the world of aviation is to secure the seat belts before start and landing, or during turbulence. In these cases, the criteria for changing from a normal state to a state of readiness are clear. In other cases it may be less obvious either because of a lack of experience or because the validity of indicators is uncertain. Examples of that can be found in the financial systems and in earthquake predictions. An obvious advantage of acting before something happens is that fewer resources may be needed, since the conditions may not yet have become critical.

All systems must be able to respond or change the way things are done when something has happened, since they otherwise will become extinct. (The only theoretically possible exception are systems for which the environment is perfectly predictable, for instance because it never changes.) The obvious advantage of proactive adjustments is that they may 'buy time,' whereas reactive adjustments always will 'take time'. Proactive adjustments can be strategic or tactical, depending on their time horizon. The potential gain is unfortunately limited by the uncertainty of whether a chosen response is the right one. On the other hand, control that is based exclusively on feedback, that is, on responding when something has happened, may quickly deteriorate into opportunistic 'fire fighting' and eventually scrambled responses, leading to a loss of control (Hollnagel, 1993).

The Four Essential Capabilities of Resilience

The above working definition of resilience can be made more concrete by considering four essential capabilities of resilience (cf., Figure E.1), namely:

- Knowing what to *do*, or being able to *respond* to regular and irregular variability, disturbances, and opportunities either by adjusting the way things are done or by activating ready-made responses. This is the capability to address the *actual*.
- Knowing what to *look for*, or being able to *monitor* that which changes, or may change, so much in the near term that it will require a response. The monitoring must cover the system's own performance as well as changes in the environment. This is the capability to address the *critical*.
- Knowing what to *expect*, or being able to *anticipate* developments, threats, and opportunities further into the future, such as potential disruptions or changing operating conditions. This is the capability to address the *potential*.
- Knowing what *has happened*, or being able to *learn* from experience, in particular to learn the right lessons from the right experience. This is the capability to address the *factual*.

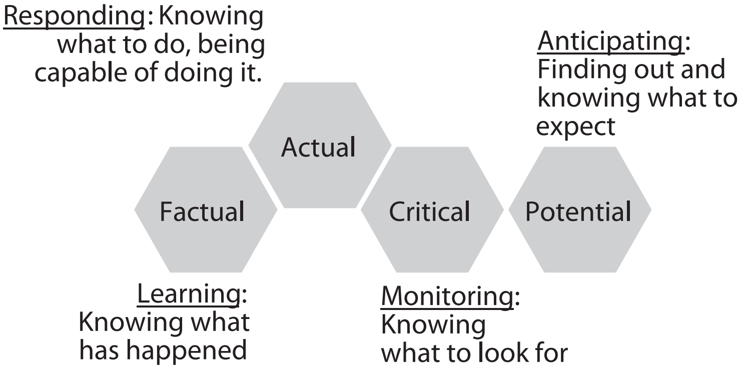
Responding: Knowing what to do, being capable of doing it.
Anticipating: Finding out and knowing what to expect
Actual
Factual
Critical
Potential
Learning: Knowing what has happened
Monitoring: Knowing what to look for

Figure E.1 The four main capabilities of a resilient system

Resilience Indicators

As mentioned in the Introduction, three things are necessary in order to be able to manage a process. The first requirement is to know what the current status or position is, in other words, to find appropriate indicators or measures – but of resilience rather than of safety. This involves several critical issues:

- Can the values of the indicators be rendered in a concise manner, either quantitative or qualitative?
- Are the indicators well defined, reliable and valid?
- Are the indicators objective, meaning that their interpretation is normative, or are they subjective, meaning that their interpretation depends on who looks at them?
- Are the indicators sufficiently sensitive to change, i.e., can the effects of a change be seen within a reasonable amount of time? (Another way of putting that is whether the indicators make concurrent control possible.)
- Are the indicators 'lagging', 'current', or 'leading', that is, do they represent a past state, the present state, or can they be interpreted as indicating a future state or development?
- Can the indicators be used as a basis for concrete actions within the operational context?
- Are the indicators easy to use ('cheap') or are they difficult to use ('costly')?

Measurements of Safety

It is quite understandable that safety indicators or safety measurements traditionally have focused on adverse outcomes, since these represent something that any system would want to avoid. Adverse outcomes also naturally attract attention both in terms of their direct effects (loss of life, property and money) and in terms of their indirect effects (disruption of functions and production, need of recovery operations, redesign, etc.).

If safety is defined by the absence of unwanted events, the level of safety is consequently measured by the relative occurrence of such events. (In fact, the definition of safety is in many cases derived from the ability to make certain measurements.) Consider, for instance, the top five HSE indicators used by the oil industry:

- (number of) fatal accidents
- total recordable injury frequency (TRIF)
- lost-time injury frequency (LTIF)
- serious HSE incident frequency (SIF)
- accidental oil spill (number and volume).

Common to these indicators are that they are reasonably objective, easy to quantify, and that they can be used without requiring costly changes to the existing system. They are probably also reliable, but it can be questioned whether they are valid safety indicators. (Another way to look at it is to ask which definition of safety the indicators imply.) They are all lagging indicators, and may be more useful to confirm effects after a while than to manage changes. Since the indicators represent outcomes rather than processes, they provide a useful basis for actions within the operational context.

This approach can be found in other industries as well. In the area of patient safety, for instance, the OECD has proposed the sets of indicators for 'operative and post-operative complications', 'sentinel events', 'hospital-acquired infections', 'obstetrics' and 'other care-related adverse events'. Here the first group contains the following indicators:

- complications of anaesthesia
- postoperative hip fracture
- postoperative pulmonary embolism (PE) or deep vein thrombosis (DVT)
- postoperative sepsis
- technical difficulty with procedure.

Similar comments can be made as for the off-shore safety indicators. The patient safety indicators all refer to well defined events, but the problem is that counting how many events there are in each category does not by itself say much about what the level of safety is.

A final example is found in the programme of work for the European Technology Platform on Industrial Safety (ETPIS, 2005). The aim of this group is to implement a new safety paradigm, called an 'incident elimination culture,' in European industry by 2020. Safety is highlighted as a key factor for successful business and as an inherent element of business performance. The aim is to demonstrate a measurable improvement of industrial safety performance by a reduction in the number of the following four categories of outcomes:

- reportable accidents at work
- occupational diseases
- environmental incidents
- accident-related production losses.

The two milestones defined by the European Technology Platform on Industrial Safety are a 25 per cent reduction in accidents by 2020 and that programmes are in place by 2020 to continue accident reduction at a rate of >5 per cent per year. While these milestones have the advantage of looking very concrete and verifiable, they also point to the main problem with commonly used safety indicators, namely that they work best in the beginning when safety is bad, but less well later when safety is good. The reason is simply that if the number of reported events is large, as it typically is when a programme of improvement is begun, then it will be easy to see a reduction in the number of adverse outcomes. But if a programme has been running successfully for some time, then there will be few reportable events to measure. This can be illustrated by Figure E.2. (The effect of a given level of effort in the beginning, Δ_1, is much larger than the effect of the same level of effort, Δ_2, later on.)

From a control or management point of view the diminishing number of outcomes is a problem, since the absence of measurements means that there is no feedback, hence that the process becomes unmanageable. The logical consequence is to look for measurements that increase rather than decrease as the situation improves.

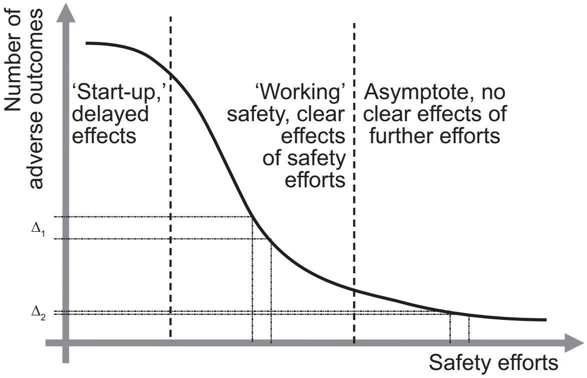

Figure E.2 The dilemma of basing safety on measuring adverse outcomes

Measurements of Resilience

Since resilience is defined by the system's ability to adjust the way things are done, it follows that a measure of resilience must be different from the traditional measures of safety. And because resilience refers to a quality rather than a quantity, to something that the system does rather than to something that the system *has*, it is highly unlikely that it can be represented by a single or simple measurement. A possible solution is instead to consider the four capabilities that together define resilience, and from that basis develop a *Resilience Analysis Grid*, that is, four sets of questions where the answers can be used to construct a resilience profile. The rest of this chapter will present a general outline of what a *Resilience Analysis Grid* (RAG) may look like.

The Ability to Respond

No system, organisation, or organism can survive unless it is able to respond to what happens – whether it is a threat or an opportunity. Responses must furthermore be both timely and effective so that they can bring about the desired outcome or change before it is too late. In order to respond, the system must first *detect* that something has happened, then recognise the event and rate it as being so serious that a response is necessary and finally know *how* and *when* to respond and be *capable* of responding.

If an event is rated as serious, the response can either be to change from a state of normal operation to a state of readiness or to take specific action in the concrete situation. In order to take action it is necessary either to have prepared responses and the requisite resources, or to be flexible enough to make the necessary resources available when needed. In responding to events, it is essential to be able to distinguish between what is *urgent* and what is *important*.

Table E.1 Probing questions for the ability to respond

	Analysis item (ability to respond)
Event list	Is there a list of events for which the system has prepared responses? Do the events on the list make sense and is the list complete?
Background	Is there a clear basis for selecting the events? Is the list based on tradition, regulatory requirements, design basis, experience, expertise, risk assessment, industry standard, etc.?
Relevance	Is the list kept up-to-date? Are there rules/guidelines for when it should be revised (e.g., regularly or when necessary?) On which basis is it revised (e.g., event statistics, accidents)?
Threshold	Are there clear criteria for activating a response? Do the criteria refer to a threshold value or a rate of change? Are the criteria absolute or do they depend on internal/external factors? Is there a trade off between safety and productivity?
Response list	How is it determined that the responses are adequate for the situations they refer to? (Empirically, or based on analyses or models?) Is it clear how the responses have been chosen?
Speed	How soon can an effective response begin? How fast can full response capability be established?
Duration	For how long can an effective response be sustained? How quickly can resources be replenished? What is the 'refractory' period?
Resources	Are there adequate resources available to respond (people, materials, competence, expertise, time, etc.)? How many are kept exclusively for the prepared responses?
Stop rule	Is there a clear criterion for returning to a 'normal' state?
Verification	Is the readiness to respond maintained? How and when is the readiness to respond verified?

The Ability to Monitor

A resilient system must be able flexibly to monitor its own performance as well as changes in the environment. Monitoring enables the system to address possible near-term threats and opportunities before they become reality. In order for the monitoring to be flexible, its basis must be assessed and revised from time to time.

Monitoring can be based on 'leading' indicators that are *bona fide* precursors for changes and events that are about to happen. The main difficulty with 'leading' indicators is that the interpretation requires an articulated description, or model, of how the system functions. In the absence of that, 'leading' indicators are defined by association or spurious correlations. Because of this, most systems rely on current and lagging indicators, such on-line process measurements and accident statistics. The dilemma of lagging indicators is that while the likelihood of success increases the smaller the lag is (because early interventions are more effective than late ones), the validity or certainty of the indicator increases the longer the lag (or sampling period) is.

Table E.2 Probing questions for the ability to monitor

	Analysis item (ability to monitor)
Indicator list	How have the indicators been defined? (By analysis, by tradition, by industry consensus, by the regulator, by international standards, etc.)
Relevance	When was the list created? How often is it revised? On which basis is it revised? Is someone responsible for maintaining the list?
Indicator type	How appropriate is the mixture of 'leading', 'current' and 'lagging' indicators? Do indicators refer to single or aggregated measurements?
Validity	For 'leading' indicators, how is their validity established? Are they based on an articulated process model?
Delay	For 'lagging' indicators, what is the duration of the lag?
Measurement type	How appropriate are the measurements? Are they qualitative or quantitative? (If quantitative, is a reasonable kind of scaling used?) Are the measurements reliable?
Measurement frequency	How often are the measurements made? (Continuously, regularly, now and then?)
Analysis / interpretation	What is the delay between measurement and analysis/interpretation? How many of the measurements are directly meaningful and how many require analysis of some kind? How are the results communicated and used?
Stability	Are the effects that are measured transient or permanent? How is this determined?
Organisational support	Is there a regular inspection scheme or schedule? Is it properly resourced?

The Ability to Anticipate

While monitoring makes immediate sense, it may be less obvious that it is useful to look at the more distant future as well. The purpose of looking at the potential is to identify possible future events, conditions, or state changes that may affect the system's ability to function either positively or negatively.

Risk assessment focuses on future threats and is suitable for systems where the principles of functioning are known, where descriptions do not contain too many details, where descriptions can be made relatively quickly and where the systems – and their environments – are sufficiently stable for their descriptions to remain valid for a reasonable time after they have been made. Many of the present day systems where industrial safety is a concern are unfortunately not like that, but are rather underspecified. For such systems the principles of functioning are only partly known, descriptions contain (too) many details, and it takes so long to make them that the system will have changed in the meantime. The systems are consequently intractable. For such systems, established risk assessment methods may be inappropriate.

The anticipation of future opportunities has little support in current methods, although it rightly ought to be considered just as important as the search for threats. This shortcoming is at least acknowledged by Resilience Engineering.

Table E.3 Probing questions for the ability to anticipate

	Analysis item (ability to anticipate)
Expertise	Is there expertise available to look into the future? Is it in-house or outsourced?
Frequency	How often are future threat and opportunities assessed? Are assessments (and re-assessments) regular or irregular?
Communication	How well are the expectations about future events communicated or shared within the organisation?
Assumptions about the future (model of future)	Does the organisation have a recognisable 'model of the future'? Is this model clearly formulated? Are the model or assumptions about the future explicit or implicit? Is the model articulated or a 'folk' model (e.g., general common sense)?

Table E.3 *Concluded*

	Analysis item (ability to anticipate)
Time horizon	How far does the organisation look ahead? Is there a common time horizon for different parts of the organisation (e.g., for business and safety)? Does the time horizon match the nature of the core business process?
Acceptability of risks	Is there an explicit recognition of risks as acceptable and unacceptable? Is the basis for this distinction clearly expressed?
Aetiology	What is the assumed nature of future threats? (What are they and how do they develop?) What is the assumed nature of future opportunities? (What are they and how do they develop?)
Culture	To which extent is risk awareness part of the organisational culture?

The Ability to Learn

It is indisputable that future performance only can be improved if something is learned from past performance. Indeed, learning is generally defined as 'a change in behaviour as a result of experience'.

The effectiveness of learning depends on the basis for learning, that is, which events or experiences are taken into account, as well as on how the events are analysed and understood.

In order for effective learning to take place there must be sufficient *opportunity* to learn, events must have some degree of similarity, and it must be possible to confirm that something has been learned. (This is why it is difficult to learn from rare events.) Learning is not just a random change in behaviour but a change that makes certain outcomes more likely and other outcomes less likely. It must therefore be possible to determine whether the learning (the change in behaviour) has the desired effect. If learning has had no effect, then it has probably not happened. And if learning has the opposite effect, then it has certainly been wrong.

In learning from experience it is important to separate what is *easy* to learn from what is *meaningful* to learn. Experience is often couched in terms of the number or frequency of occurrence of adverse events. But compiling extensive accident statistics does not mean that anyone will actually learn anything. Furthermore, since the number of things that go right, including near misses, is many orders of magnitudes larger than the number of things that go wrong, it makes good sense to try to learn from representative events rather than from failures alone.

Table E.4 Probing questions for the ability to learn

	Analysis item (ability to learn)
Selection criteria	Is there a clear principle for which events are investigated and which are not (severity, value, etc.)? Is the selection made systematically or haphazardly? Does the selection depend on the conditions (time, resources)?
Learning basis	Does the organisation try to learn from what is common (successes, things that go right) as well as from what is rare (failures, things that go wrong)?
Data collection	Is there any formal training or organisational support for data collection, analysis and learning?
Classification	How are the events described? How are data collected and categorised? Does the categorisation depend on investigation outcomes?
Frequency	Is learning a continuous or discrete (event-driven) activity?
Resources	Are adequate resources allocated to investigation/analysis and to dissemination of results and learning? Is the allocation stable or is it made on an *ad hoc* basis?
Delay	What is the delay between the reporting the event, analysis, and learning? How fast are the outcomes communicated inside and outside of the organisation?
Learning target	On which level does the learning take effect (individual, collective, organisational)? Is there someone responsible for compiling the experiences and making them 'learnable'?
Implementation	How are 'lessons learned' implemented? Through regulations, procedures, norms, training, instructions, redesign, reorganisation, etc.?
Verification/ maintenance	Are there means in place to verify or confirm that the intended learning has taken place? Are there means in place to maintain what has been learned?

Applying the RAG – Rating Resilience

By considering in detail each of the four capabilities that define resilience, it is possible to propose four sets of issues and four corresponding sets of questions that can serve as a basis for assessing a system's resilience. The same set of issues can also be the starting point for possible concrete measures to maintain or improve resilience.

The four sets of issues together comprise what is called the *Resilience Analysis Grid* (RAG). Similarly, the answers to the four sets of questions characterise the resilience of a system and can be used to construct a *resilience profile*. It is, of course, possible to work just with the answers or ratings, but for many purposes it is useful also to have some kind of pictorial or graphical representation to help communicate and discuss the results. The so-called *star chart* or radar chart is well suited for this purpose. The star chart is a straightforward way to display multivariate data in the form of a two-dimensional chart where all the variables are represented on axes starting from the same point, cf. Figure E.3.

The procedure for filling out a RAG is quite simple, and can be described the following steps.

Define and Describe the System for which the RAG is to be Constructed

The first step is, not surprisingly, to provide a clear and concise description of the system for which the RAG is to be filled out. Is the system, for instance, an aircraft crew (pilots plus flight attendants), the flight dispatch service, aircraft maintenance, or the airline as a whole? Is the system the central control room of a power plant, a work shift, the maintenance and repair services or the outage handling? A resilience analysis must always begin by defining as clearly as possible the boundaries of the system being considered, the organisational structure, the people and resources involved, the time horizon for typical activities, etc. Without that it is not possible to know which questions to ask, nor how to rate the answers.

Select a Subset of Relevant Questions for Each of the Four Capabilities

The four sets of questions presented in this chapter do not refer to any specific system or domain and should therefore not be used without confirming their relevance. This can be done in two steps. The first is to select four subsets of questions that correspond to the system defined by the first step, that is, the scope of activities and the nature of the core processes. The second is to reformulate individual questions so that they are appropriate for the domain, and possibly add new questions if needed. An investigation of the resilience of a hospital ward should, for instance, not use the same questions and the same formulations as an off-shore drilling rig. Different domains, and different kinds of businesses, may also affect the relative weight or importance of the four capabilities.

Rate the Selected Questions for Each Capability

Based on the outcome of the second step, it is now possible to get answers to the four sets of questions and to rate the answers. The answers must come from people who have experience of the domain. Various approaches can be used such as workplace interviews, discussions with experts, focus groups, etc. Since it is important for the proper use of the RAG that the ratings are done repeatedly rather than only once, it may be useful to nominate a number of people in the system who can serve as a pool of respondents.

In order for the RAG to be useful as a tool, it is necessary that the answers to each question are rated using a common terminology. It is proposed to use the following five categories.

- *Excellent* – the system on the whole *exceeds* the criteria addressed by the specific item.
- *Satisfactory* – the system fully *meets* all *reasonable* criteria addressed by the specific item.
- *Acceptable* – the system meets the *nominal* criteria addressed by the specific item.
- *Unacceptable* – the system does *not* meet the *nominal* criteria addressed by the specific item.
- *Deficient* – there is *insufficient* capability to meet the criteria addressed by the specific item.

In addition, a sixth category must be included to account for the situation where the system does not address a capability at all.

- *Missing* – there is *no* capability whatsoever to address the specific item.

When answering and rating the individual items in the lists, it should be kept in mind that the rating is not intended to be a 'scoring' of recent accidents and incidents. The examples that are used to provide the answers should be of the normal or typical way in which the system functions. If there has been a number of cases where the system has failed to meet the criteria, then this should clearly been taken into account during the rating. But the rating should describe how well the system is able to do something, rather than how badly things can turn out.

Combine the Ratings to a Score for Each Capability, and for the Four Capabilities Combined

Once the rating has been done for each set of items, they can be shown by means of a star chart. To illustrate this, Figure E.3 shows an empty star chart for the ability to monitor. The star chart has ten axes, corresponding to the ten variables (items) used to rate the ability to monitor, with each axis marked using the five rating categories described above. The sixth category of *missing* corresponds to the common starting point of the axes.

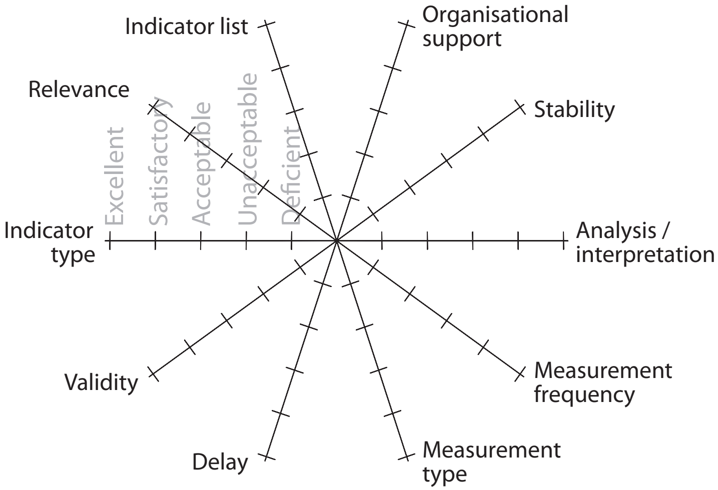

Figure E.3 Empty star chart for monitoring

The star chart is used in the following way. If, for instance, all variables were rated as 'acceptable', then the result would be a regular polygon (not shown in Figure E.3). If one or more of the variables were rated differently, either better or worse, then the result would be an irregular polygon. The shape of the polygon that is constructed from the ratings therefore provides a convenient visual representation of the 'balance' among the ratings. Note, however, that the reference rating for a specific variable will depend on the nature of the system's activities. A specific system may, for instance, require that the 'validity' is *excellent*, whereas the 'measurement type' (i.e., mixture of qualitative and quantitative measures) only has to be *acceptable*. The star charts for the other abilities are produced in the same straightforward manner. The star charts for the four capabilities will together provide an overall view of how the system's resilience was rated.

It is, however, also possible to combine the four star charts into one by making comparable the several dimensions (axes) for each capability. The simplest approach is to assign numerical values to the ratings, for instance from 1 to 5 where 1 corresponds to 'deficient,' 2 to 'unacceptable', and so on. It is then straightforward to calculate the value of the rating for each axis and to aggregate them into a single value. This approach can be made more reasonable by assigning appropriate weights to both the ratings and the dimensions. Provided that a procedure can be defined that respects the characteristics of the system and the domain, the RAG can be represented by a four-axis star diagram, as shown in Figure E.4.

The assignments shown in Figure E.4 are for purpose of illustration only. In this example the shape of the polygon is irregular, indicating that all is not well. The figure corresponds to a system that does well in terms of the ability to respond and monitor, but which fails in terms of the ability to anticipate and learn. While such a system may be safe in the short run, it is not resilient.

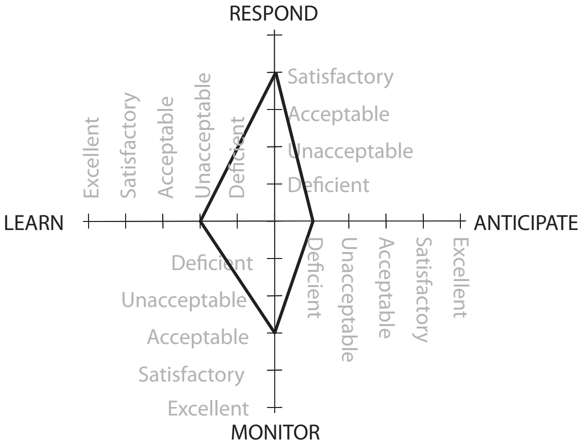

Figure E.4 An aggregated star diagram

Interdependence of the Four Capabilities

A less simple but more meaningful approach is to consider how the four basic capabilities are coupled. The ability to respond, for instance, can be enhanced by monitoring, which in turn may benefit from learning. While a detailed description of the couplings is beyond the scope of this chapter, a first attempt could be as follows (cf., Figure E.5), using the FRAM representation described in Chapter 13.

Responses can be triggered by external and/or internal events, and this can be facilitated by the output from the monitoring function. The response itself requires that the system is in a state of readiness and that the necessary resources (tools, materials and people) are available. The scheduling of the response is controlled by plans and procedures, predefined or *ad hoc*, and may require that the scheduling of ongoing actions is flexible so that the normal activities can be resumed when the response has come to an end.

The input to *monitoring* comes from internal and external developments that provide the raw data, and from the functions of anticipation and learning that provide the background for looking at and interpreting the data. Effective *monitoring* requires both that there is time available (cf., Hollnagel, 2009), that there is a monitoring strategy (i.e., that monitoring is both efficient and thorough), and that the people or operators involved have the requisite skills and knowledge.

Anticipation is heavily influenced by what has been learned from the past, such as suggestions for performance indicators. It is controlled or guided by the 'model of the future'' in particular the types of threats or opportunities that this model describes. Unlike the other functions, anticipation is not necessarily data-driven. The main resource is competent people, but anticipation is rarely a time-critical function. The pre-condition is the organisational culture or awareness, here described as a 'constant sense of unease', cf., Hollnagel et al. (2008).

Learning, finally, makes use of past events and responses, either in-house or in the general domain of activity, possibly mediated by regulators, and internal or external events even if these have not resulted in something requiring a response. Learning is 'controlled' or guided by the assumptions about why things happen. Here the organisation's accident model is of particular importance, for instance in the way in which it determines which data and events are considered (Lundberg et al., 2009). Effective learning finally requires some kind of reporting scheme.

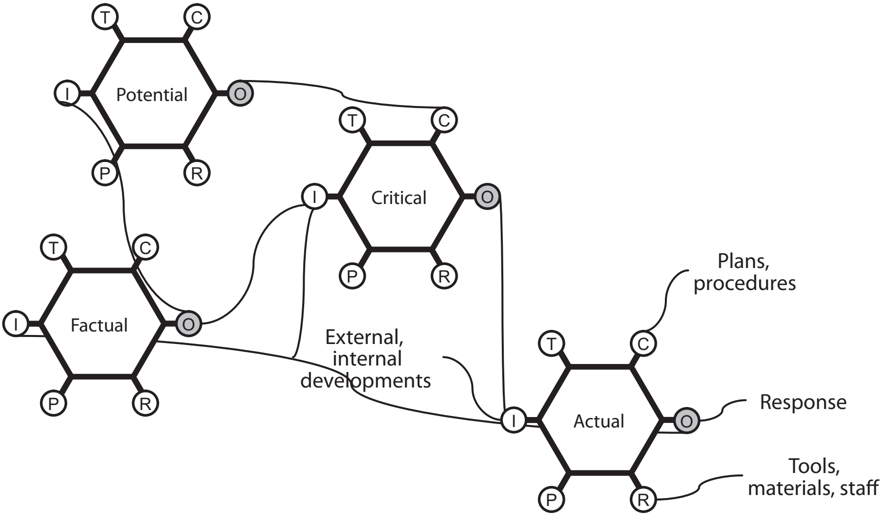

Figure E.5 Interdependence of the four resilience capabilities

Single versus Repeated Measures

Since the RAG is intended as a tool to support resilience management rather than resilience measurement, it is essential that it is used regularly. The RAG should not just give a measurement of a system's resilience at a single point in time, but be used to follow how resilience develops over time. The RAG is thus itself intended as a (composite) current indicator, rather than a simple lagging or a leading indicator. When used in this way it actually becomes less critical how the aggregated star chart is produced, since the relative indications are more important than the absolute. But it is important that the RAG is applied systematically and consistently.

The frequency of the ratings clearly depends on the characteristics of the system's core business and on the volatility of the operating environment. It is therefore not possible to provide any strict guidelines for that. But given the dynamics of current societies it does seem sensible to perform a rating for a system or an organisation at least every 2–3 months. (In business it does seem to be a tradition, if not a demand, to produce a report on how well things are going four times a year.)

Summary

The *Resilience Analysis Grid* presented here shows how it is possible to develop a tool that can support resilience management. It is not a tool that can be used off-the-shelf. It is rather intended as a basis from which more specific grids – or set of questions – can be developed.

The chapter has presented and discussed the principles for how the dimensions can be rated, and how they can be shown by means of a star diagram. The star diagram is not in itself a measure of resilience, but a compact representation of how the various items were rated. The RAG should also be thought of as a process measure rather than a product measure, since it shows the current level of resilience and of how well the system does on each of the four main capabilities.

Resilience Engineering cannot prescribe a certain balance or proportion among the four qualities. For a fire brigade, for instance, it is more important to be able to respond to the actual than to consider the potential, whereas for a sales organisation, the ability to anticipate may be just as important as the ability to respond. But it is clearly necessary for any system to address each of these qualities to some extent in order to be resilient. All systems traditionally put some effort into the ability to respond to the actual. Many also put some effort into the ability to learn from the factual, although it often is in a very stereotypical manner. Fewer systems make a sustained effort to monitor the critical, particularly if there has been a long period of stability. And very few systems put any serious effort into the ability to anticipate the potential.

Bibliography

Aase, K., Wiig, S. & Høyland, S. (2009). Safety first!? Organizational efficiency trends and their influence on safety. *Safety Science Monitor*, 13(2) article 7. [12]

ACCOMPLI: http://www.aviation-civile.gouv.fr/publications.htm. [2]

Åkerstedt, T. & Gillberg, M. (1990). Subjective and objective sleepiness in the active individual. *The International Journal of Neurosciences*, 52(1–2), 29–37. [6]

Alderson, D.L. & Doyle, J.C. (2010). Contrasting views of complexity and their implications for network-centric infrastructures. *IEEE Systems, Man and Cybernetics*, Part A. 40(4), 839–52. [10]

Amalberti, R. (1996). *La Conduite des Systèmes à Risques*. Paris: PUF, Coll. Le travail humain. [3]

Amalberti, R. (2001). The paradoxes of almost totally safe transportation systems. *Safety Science,* 37(2–3), 109–126. [3] [18]

Amalberti, R. (2006). Optimum system safety and optimum system resilience: Agonistic or antagonistic concepts? In E. Hollnagel, D.D. Woods & N.G. Leveson (eds), *Resilience Engineering: Concepts and Precepts*. Aldershot, UK: Ashgate (pp. 253–74). [Prologue] [18]

Amalberti, R., Auroy, Y., Berwick, D. & Barach, P. (2005). Five system barriers to achieving ultrasafe health care. *Annals of Internal Medicine*, 142, 756–64. [3]

Andersson, K.P. & Ostrom, E. (2008). Analyzing decentralized resource regimes form a polycentric perspective. *Policy Science*, 41, 71–93. [9]

Babin, L., Lee, S., Halko, S., Boudreau, A.C. & George, C.F.P. (1997). Determining sleep-wake activity using Actiwatch. *Sleep Research*, 26, 640. [6]

Bærentsen, K.B. (1996). Episodic knowledge in system control. In B. Holmqvist, P.B. Andersen, H. Klein & R. Posner (eds), *Signs of Work: Semiosis and Information.* Berlin: Walter de Gruyter & Co (pp. 283–319). [3]

Bainbridge, L. (1987). The ironies of automation. In J. Rasmussen, K. Duncan & J. Leplat (eds), *New Technology and Human Error*. London: Wiley (pp. 271–83). [2]

Bank of England (2007). *Financial Stability Report*, October 2007, Issue 22. www.bankofengland.co.uk, accessed September 2008. [13]

Bengtsson, J., Angelstam, P., Elmqvist, T., Emanuelsson, U., Folke, C., Ihse, M., Moberg, F. & Nyström, M. (2003). Reserves, resilience and dynamic landscapes. *Ambio*, 32(6), 389–96. [10]

Bergström, J., Dahlström, N., van Winsen, R., Lützhöft, M., Dekker, S.W.A. & Nyce, J. (2009). Rule- and role-retreat: An empirical study of procedures and resilience. *Journal of Maritime Research*, 6(1), 75–90. [4]

Bernstein, P.L. (2007). To botch an economic forecast, rely on past experience. *International Herald Tribune*, Sunday, December 30. [13]

Bertrand, M. & Mullainathan, S. (2001). Do people mean what they say? Implications for Subjective Survey Data. *The American Economic Review*, 91, 67–72. [18]

Bertuglia, C.S. & Vaio, F. (2005). *Non-Linearity, Chaos and Complexity* (2nd Edition). Oxford, UK: Oxford University Press. [15]

Bird Strike Committee-USA (2009). Significant bird and other wildlife strikes. http://www.birdstrike.org/commlink/signif.htm, accessed December 2009. [2]

Bird, F.E. Jr., Germain, G.L. & Clark, M.D. (2003). *Practical Loss Control Leadership* (3rd ed.). Det Norske Veritas (USA) [7]

Bisseret, A., Sebillotte, S. & Falzon, P. (1999). *Techniques Pratiques pour L'étude des Activités Expertes*. Toulouse: Octarès-Editions. [3]

Bjørnskau. T. (2005). Aviation safety in Norway: Results from a questionnaire survey to employees in Norwegian aviation (in Norwegian). The Institute of Transport Economics (TØI) report no 782/2005. [12]

Bonabeau, E. & Teraulaz, G. (1994). *Intelligence Collective*. Paris: Hermès. [16]

Bowlby, J. (1973). *Attachment and Loss*. Vol 2: Separation. New York: Basic Books. [16]

Branlat, M., Fern, L., Voshell, M. & Trent, S. (2009). Coordination in urban fire-fighting: A study of critical incident reports. *Proceedings of the Human Factors and Ergonomics Society 53rd Annual Meeting*, San Antonio, TX. [10]

Brehmer, B. (2008). *Från Funktioner till Konkret Ledningssystem för Komplexa Operationer*. Stockholm: Försvarshögskolan. [4]

Bressolle, M.C., Decortis, F., Pavard, B. & Salembier, P. (1996). Traitement cognitif et organisationnel des micro-incidents dans le domaine du contrôle aérien: analyse des boucles de régulation formelles et informelles. In G. De Terssac & E. Friedberg (eds), *Coopération et Conception*. Toulouse: Octares (pp. 267–88). [16]

Brown, J.P. (2005). Key themes in healthcare safety dilemmas. In M.S. Patankar, J.P. Brown, & M.D. Treadwell (eds), *Safety Ethics: Cases From Aviation, Healthcare, and Occupational and Environmental Health*. Aldershot, UK: Ashgate (pp. 103–48). [10]

Burke, K. (1969). *A Grammar of Motives*. Berkeley & Los Angeles, CA: University of California Press. [17]

Burr, V. (2003). *Social Constructionism*. East Sussex, UK: Routledge. [18]

Cabon, P. (2008). *De la Gestion de la Fatigue à la Gestion Organisationnelle du Risque Fatigue. Habilitation à Diriger des Recherches*. Paris, Université René Descartes. [6]

Cabon, P., Mollard, R., Debouck, F., Chaudron, L., Grau, J.Y. & Deharvengt, S. (2008). From flight time limitations to fatigue risk management systems. *3rd Symposium on Resilience Engineering*, Antibes Juan-Les-Pins, France, October, 28–30. [6]

Carruthers, I. & Philip, P. (2006). Safety first – A report for patients, clinicians and healthcare managers. Department of Health. http://www.dh.gov.uk/prod_consum_dh/groups/dh_digitalassets/@dh/@en/documents/digitalasset/dh_064159.pdf, accessed June 2010. [17]

Carthey, J., De Leval, M.R., Wright, D.J., Farewell, V.T. & Reason, J.T. (2003). Behavioural markers of surgical excellence. *Safety Science*, 41(5), 409–25. [2]

Civil Aviation Authority (1998). *CAP 681: Global Fatal Accident Review 1980–1996*. http://www.caa.co.uk/docs/33/CAP681.pdf, accessed January 2008. [18]

Civil Aviation Authority of New Zealand (2007). *Part 121 Air Operations: Large Aeroplanes*. Wellington: Civil Aviation Authority of New Zealand. [6]

Cohen, M.S., Freeman, J.T. & Wolf, S.P. (1996). Meta-recognition in time stressed decision making: Recognizing critiquing and correcting. *Human Factors*, 38, 206–19. [8]

Committee on the Future of Emergency Care in the US (2006). *Hospital-Based Emergency Care: At the Breaking Point*. Washington, DC: National Academic Press. [10]

Committee on Transportation and Infrastructure, (2009). http://transportation.house.gov/hearings/Testimony.aspx?TID=12805&NewsID=809, accessed July, 2010. [2]

Cook, R.I. (2006). Being bumpable: Consequences of resource saturation and near-saturation for cognitive demands on ICU practitioners. In D.D. Woods & E. Hollnagel (eds), *Joint Cognitive Systems: Patterns in Cognitive Systems Engineering*. Boca Raton, FL: CRC Press. Taylor & Francis Group (pp. 23–35). [10]

Cook, R.I. & Nemeth, C.P. (2006). Taking things in one's stride: Cognitive features of two resilient performances. In E. Hollnagel, D.D. Woods & N.G. Leveson (eds), *Resilience Engineering: Concepts and Precepts*. Aldershot, UK: Ashgate (pp. 205–21). [Epilogue]

Cook, R.I. & Rasmussen, J. (2005). "Going solid": A model of system dynamics and consequences for patient safety. *Quality and Safety in Health Care*, 14, 130–34. [9]

Cook, R.I. & Woods, D.D. (2006). Distancing through differencing: An obstacle to learning following accidents. In E. Hollnagel, D. D. Woods and N. Leveson (eds), *Resilience Engineering: Concepts and Precepts*. Aldershot, UK: Ashgate. (pp. 329–38.) [10]

Cook, R.I., Woods, D.D. & McDonald, J.S. (1991). *Human Performance in Anesthesia: A Corpus of Cases*. Cognitive Systems Engineering Laboratory Report, prepared for Anesthesia Patient Safety Foundation, April 1991. [10]

Cook, R.I., Woods, D.D. & Miller, C. (1998). *A Tale of Two Stories: Contrasting Views of Patient Safety*. North Adams, MA: US National Patient Safety Foundation. [10] [17]

Coutarel, F., Daniellou, F. & Dugué, B. (2003). Interroger l'organisation du travail au regard des marges de manoeuvre en conception et en fonctionnement [Examining work organization in relation to margins of manoeuvre in design and in operation]. *Pistes*, 5(2). Online at http://www.pistes.uqam.ca/v5n2/articles/v5n2a2.htm, accessed April 2010. [10]

Csete, M.E. & Doyle, J.C. (2002). Reverse engineering of biological complexity. *Science*, 295, 1664–1669. [10]

Cuvelier, L. & Falzon, P. (2008). Methodological issues in the quest of resilience factors. In E. Hollnagel, F. Pieri & E. Rigaud (eds), *3rd International Symposium on Resilience Engineering*. October 28–30, Antibes – Juan-les-Pins, France. [3]

Cyrulnik, B. (2003). *Le Murmure des Fantômes*. Paris: Odile Jacobs. [16]

De Keyser, V. & Nyssen, A.S. (1993). L'erreur humaine en anesthésie. *Le travail humain*, 56, 243–66. [16]

de Larosière, J. (2009). *The High-Level Group of Financial Supervision in the EU*. Brussels, February 25. http://ec.europa.eu/internal_market/finances/docs/de_larosiere_report_en.pdf, accessed February 2010. [13]

Deharvengt, S. (2007). Barriers to Regulating Resilience: Example of Pilots' Crew Resource Management Training. Paper presented at the Resilience Engineering Workshop, Vadstena, Sweden, 25–27 June. [6]

Dekker, S.W.A. (1996). Cognitive complexity in management by exception: deriving early human factors requirements for an envisioned air traffic management world. In D. Harris (ed.), *Engineering Psychology and Cognitive Ergonomics.* Vol. I. Aldershot, UK: Ashgate (pp. 201–210). [8]

Dekker, S.W.A. (2005). *Ten Questions About Human Error: A New View of Human Factors and System Safety*. New Jersey, USA: Lawrence Erlbaum Publishers. [18]

Dekker, S.W.A. (2006). Resilience engineering: Chronicling the emergence of confused consensus. In E. Hollnagel, D.D. Woods, and N. Leveson (eds), *Resilience Engineering Concepts and Precepts*. Farnham, UK: Ashgate (pp. 77–92). [12]

Dekker, S.W.A. (2007). *Just Culture: Balancing Safety and Accountability*. Aldershot, UK: Ashgate. [10] [17]

Dekker, S.W.A., Dahlström, N., van Winsen, R. & Nyce, J. (2008). Crew resilience and simulator training in aviation. In E. Hollnagel, C. Nemeth & S.W.A. Dekker (eds), *Resilience Engineering Perspectives, Remaining Sensitive to the Possibility of Failure.* Aldershot, UK: Ashgate (pp. 119–26). [4]

Dekker, S.W.A. & Woods, D.D. (1999). To intervene or not to intervene: The dilemma of management by exception. *Cognition Technology and Work*, 1, 86–96. [8]

Dekker, S.W.A. & Woods, D.D. (2002). MABA-MABA or abracadabra? Progress on human-automation co-ordination. *Cognition, Technology & Work*, 4, 240–4. [18]

Department of Health. (2000). *An Organisation With a Memory: Learning From Adverse Events in the NHS*. London, UK: The Stationary Office. [17]

Dörner, D. (1996). *The Logic of Failure.* New York: Metropolitan Books. [4]

Downer, J. (2009). Epistemological Chicken: What Do We Learn From Aircraft 'Bird-Ingestion' Tests? Centre for Analysis of Risk and Regulation, The Economic and Social Research Council (ESRC), London School of Economics and Political Science, Houghton Street, London, UK. [2]

Doyle, J.C. (2000). Multiscale networking, robustness, and rigor. In T. Samad & J. Weyrauch (eds), *Automation, Control, and Complexity: An Integrated Approach*. New York: John Wiley & Sons, Inc. (pp. 287–301). [1] [10]

ESSAI: http://essai.nlr.nl/introduction.htm. [2]

EUROCONTROL. (2000). ESARR 3 – EUROCONTROL Safety Regulatory Requirement. Use of Safety Management Systems by ATM Service Providers. [17]

European Technology Platform on Industrial Safety (ETPIS) (2005). *Safety for Sustainable European Industry Growth: Strategic Research Agenda*. www.industrialsafety-tp.org accessed July 2010. [Epilogue]

Falzon, P. (2005). Ergonomics, knowledge development and the design of enabling environments. In *Humanizing Work and Work Environment Conference*. December 10–12, Guwahati, India. [3]

Federal Aviation Administration (FAA). (2009). Transcript of New York TRACON La Guardia Departure Air Traffic Controller communications. [2]

Feltovich, P.J., Spiro, R.J. & Coulson, R.L. (1989). The nature of conceptual understanding in biomedicine: The deep structure of complex ideas and the development of misconceptions. In D. Evans & V. Patel (eds), *The Cognitive Sciences in Medicine*. Cambridge MA: MIT Press (pp. 113–72). [10]

Feltovich, P.J., Spiro, R.J. & Coulson, R.L. (1997). Issues of expert flexibility in contexts characterized by complexity and change. In P.J. Feltovich, K.M. Ford, & R.R. Hoffman (eds), *Expertise in Context: Human and Machine*. Menlo Park, CA: AAAI/MIT Press (pp. 113–172). [10]

Ferguson, E. & Cox, T. (1993). Exploratory factor analysis: A users' guide. *International Journal of Selection and Assessment*, 1(2), 84–94. [11]

Ferreira, P., Clarke, T., Wilson, J.R., Sharples, S. & Ryan, B. (2008). Resilience in Rail Engineering Work. In E. Hollnagel, F. Pieri & E. Rigaud (eds), *Proceedings of the 3rd Resilience Engineering Symposium* (October, 28–30, 2008 Antibes, France). Paris, France: Ecole des Mines de Paris. [11]

Financial Stability Board (April, 2009). *Report of the Financial Stability Forum on Addressing Procyclicality in the Financial System*. http://www.financialstabilityboard.org/publications/r_0904a.pdf, accessed February 2010. [13]

Financial Stability Forum (2008). *FSF Working Group on Market and Institutional Resilience – Interim Report*. February 2008. http://www.financialstabilityboard.org/publications/r_0802.htm, accessed February 2010. [13]

Financial Stability Forum (2008). *Report of the Financial Stability Forum on Enhancing Market and Institutional Resilience*. April 2008. http://www.financialstabilityboard.org/publications/r_0804.pdf, accessed February 2010. [13]

Flanagan, J.C. (1954). La technique de l'incident critique. *Revue de Psychologie Appliquée*, 4(3), 267–95. [3]

Fletcher, G., Flin, R., McGeorge, P., Glavin, R., Maran, N. & Patey R. (2004). Rating non-technical skills: Developing a behavioural marker system for use in anaesthesia. *Cognition Technology and Work*, 6, 165–71. [8]

Flin, R., Martin, L., Goeters, K., Hoermann, J., Amalberti, R., Valot, C. & Nijhuis, H. (2003). Development of the NOTECHS (Non-Technical Skills) system for assessing pilots' CRM skills. *Human Factors and Aerospace Safety*, 3, 95–117. [8]

Folkard, S. & Åkerstedt, T. (2004). Trends in the risk of accidents and injuries and their implications for models of fatigue and performance. *Aviation, Space and Environmental Medicine*, 75(3, Suppl.): A161–7. [6]

Foushee, C.H., Lauber, J.K., Baetge, M.M. & Acomb, D.B. (1986). Crew Factors in flight operations. Part 3: The operational significance of exposure to short-haul air transport operations. Moffett Field, CA Ames Research Center, National Aeronautics and Space Administration: 67. [6]

Fowlkes, J., Dwyer, D., Oser, R., (1998). Event-based approach to training. *International Journal of Aviation Psychology*, 8, 209–22. [8]

Freer, D. (1994). ICAO at 50 years: Riding the flywheel of technology. *ICAO Journal*, 49(7), 9 19–32. [15]

Gaba, D. M. (2000). Structural and organizational issues in patient safety. A comparison of health care to other high-hazard industries. *California Management Review*, 43(1), 83–102. [12]

Gallati, R. (2003). *Risk Management and Capital Adequacy*. New York, NY: McGraw Hill. [13]

Gergen, K.J. (1999). *An Invitation to Social Constructionism*. London: UK: Sage Publications Ltd. [18]

Gilbert, N. & Conte, R. (1995). *Artificial Societies: The Computer Simulation of Social Life*. London, UCL Press [16].

Grote, G. (2004). Uncertainty management at core of the system design. *Annual Reviews in Control*, 28(2), 267–74. [12]

Grote, G. (2008). Rules management as source for loose coupling in high-risk systems. In E. Hollnagel, C.P. Nemeth & S.W.A. Dekker (eds), *Remaining Sensitive to the Possibility of Failure, Resilience Engineering Perspectives, Volume 1*. Farnham, UK: Ashgate (pp. 91–100). [12]

Guérin, F., Laville, A., Daniellou, F., Duraffourg, J. & Kerguelen, A. (1997). *Comprendre le travail pour le transformer. La pratique de l'ergonomie*. Lyon. ANACT [3]

Hale, A.R. (2009). Why safety performance indicators? *Safety Science*, 47(4), 479–80. [5]

Hale, A.R. & Swuste, P. (1998). Safety rules: procedural freedom or action constraint. *Safety Science*, 29, 163–77. [12]

Hauland, G., Serck-Hanssen, C. & Rolfsen, J. (2007). Exploring methodology for change processes: An aviation case of combined behaviour- and culture change to improve safety. In T. Aven and J. E Vinnem (eds), *Risk Reliability and Societal Safety*, *Volume 2*, 1665–1662. Taylor & Francis. [12]

Heath, R. (1998). Dealing with the complete crisis-the crisis management shell structure. *Safety Science*, 30(1–2), 139–50. [4]

Helmreich, R.L. & Merritt, A.C. (1998). *Culture at Work in Aviation and Medicine.* Aldershot, UK: Ashgate. [18]

Hendriksen, B. (2008). *Feasibility of the Complexity Theory in Learning from Naval Disasters*. Masters Thesis Faculty of Policy, Technology and Management. Delft University of Technology, [15]

Herrera, I.A., Norsdkag, A.O., Myhre, G. & Halvorsen, K. (2009). Aviation safety and maintenance under major organizational changes, investigating non-existing accidents. *Accident Analysis and Prevention*, 41(6), 1155–63. [14]

Hofstede, G. (2001). *Culture's Consequences: Comparing Values, Behaviors, Institutions, and Organizations Across Nations* (2nd edition). California, USA: Sage Publications, Inc. [18]

Holling, C.S. (1973). Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics*, 4, 1–23. [18]

Hollnagel, E. (1983). Human error. Position Paper for *NATO Conference on Human Error*. Bellagio, Italy. [18]

Hollnagel, E. (1993). *Human Reliability Analysis: Context and Control*. London: Academic Press. [Epilogue]

Hollnagel, E. (2002). Understanding accidents – from root causes to performance variability. In J.J Persensky, B. Halbert & H. Blackman (eds). *Proceedings from IEEE 7th Conference on Human Factors and Power Plants. New century, New trends*. 15–19 September 2002, Scottsdale. [12]

Hollnagel, E. (2004). *Barrier and Accident Prevention.* Aldershot, UK: Ashgate. [3] [12] [13] [17] [18]

Hollnagel, E. (2007). Resilience Engineering: Why, what and how? In *NoFS 2007 – Nordic Research Conference on Safety*, 13–15 June 2007, Tampere, Finland. [10]

Hollnagel, E. (2008). From protection to resilience: Changing views on how to achieve safety. *8th International Symposium of the Australian Aviation Psychology Association*. 8–11 April, Sydney, Australia.

Hollnagel, E. (2009a). Extending the scope of the human factor. In E. Hollnagel (ed.), *Safer Complex Industrial Environments*. Boca Raton, FL: CRC Press. (pp. 37–60). [Prologue]

Hollnagel, E. (2009b). *The ETTO Principle: Efficiency-Thoroughness Trade-Off ,Why Things That Go Right Sometimes Go Wrong*. Farnham, UK: Ashgate. [10] [11] [12] [Epilogue]

Hollnagel, E., Nemeth, C.P. & Dekker, S.W.A. (2008). *Remaining Sensitive to the Possibility of Failure*. Aldershot, UK: Ashgate. [15] [Epilogue]

Hollnagel, E., Pruchnicki, S., Woltjer, R. & Etcher, S. (2008). A functional resonance accident analysis of Comair flight 5191. Paper presented at the *8th International Symposium of the Australian Aviation Psychology Association*. Sydney, Australia. [13]

Hollnagel, E. & Sundström, G.A. (2006). States of resilience. In E. Hollnagel, D.D. Woods & N. Leveson (eds), *Resilience Engineering: Concepts and Precepts*. Aldershot, UK: Ashgate (pp. 339–344). [3]

Hollnagel, E. & Woods, D.D. (1983). Cognitive systems engineering. New wine in new bottles. *International Journal of Man-Machine Studies*, 18, 583–600. [18]

Hollnagel, E. & Woods, D.D. (2005). *Joint Cognitive Systems: Foundations of Cognitive Systems Engineering*. Boca Raton, FL: CRC Press. Taylor & Francis Group. [8]

Hollnagel, E. & Woods, D.D. (2006). Epilogue: Resilience Engineering precepts. In E. Hollnagel, D. D. Woods & N. Leveson (eds), *Resilience Engineering: Concepts and Precepts*. Aldershot, UK: Ashgate. [8]

Hollnagel, E., Woods, D.D. & Leveson N.G. (2006). *Resilience Engineering: Concepts and Precepts*. Aldershot, UK: Ashgate Publishing. [2] [6] [11] [17] [18]

Hoogendorp, G. (2007). *The Bijlmermeer Aviation Disaster*. Thesis Risk Management, Faculty Aerospace Engineering, Delft University of Technology. [15]

Hopkins, A. (2009a). Reply to comments. *Safety Science*, 47(4), 508–10. [5]

Hopkins, A. (2009b). Thinking about process safety indicators. *Safety Science*, 47(4), 460–5. [5]

House, R.J., Hanges, P.J., Javidan, M., Dorfman, P.W. & Gupta, V. (2004). *Culture, Leadership, and Organizations: The GLOBE Study of 62 Societies.* California, USA: Sage Publications, Inc. [18]

Houtman, H. (2008). *Your Life Will Never be the Same After an Accident Has Happened.* Presentation at School of Aviation, Lund University Sweden. [15]

Høyland, S. & Aase, K. (2009). Does change challenge safety? Complexity in the civil aviation transport system. In S. Martorell, et al. (eds), *Safety, Reliability and Risk Analysis: Theory, Methods, and Applications*. Boca Raton, FL: CRC Press (pp. 1385–93). [12]

Høyland, S., Aase, K., Pettersen K.A. & Tjørhom, B. (2008). Risk challenges and parallel change processes within the Norwegian transportation sector (in Norwegian). *Report from University of Stavanger*, No.14. [12]

Hubbard, D.W. (2009). *The Failure of Risk Management*. USA: John Wiley & Son Inc. [12]

Hutchins, E. (1995). *Cognition in the Wild*. Cambridge, MA: MIT Press. [8]

Hutchins, E., Holder, B.E. & Pérez, R.A. (2002). *Culture and Flight Deck Operations*. White paper for The Boeing Company. Retrieved July 2008 from http://hci.ucsd.edu/lab/hci_papers/EH2002-2.pdf. [18]

ICAO (2006). *Safety Management Manual* (Doc 9859, AN/460). Montreal, Canada: International Civil Aviation Organization. [Epilogue]

ICAO (2007). *Air Traffic Management. Procedures for Air Navigation Services*, DOC 4444 ATM/501, (15th Edition). ICAO. [8]

ICAO (2008). Fatigue Risk Management Systems, *9th Meeting of the Operations Panel (OPSP)*. Montréal. [6]

Jentsch, F. & Bowers, C.A. (1998). Evidence for the validity of low-fidelity simulations in air crew coordination research and training. *International Journal of Aviation Psychology*, 8, 243–60. [8]

Johansson, B. & Hollnagel, E. (2007). Pre-requisites for large scale coordination. *Cognition, Technology & Work*, 9, 5–13. [16]

Johnson, C.W. (2002). Reasons for the Failure of Incident Reporting in the Healthcare and Rail Industries. Paper presented at the *Components of System Safety: Proceedings of the 10th Safety-Critical Systems Symposium*, Southampton, UK. [17]

Johnson, C.W. (2003). *Failure in Safety-Critical Systems: A Handbook of Incident and Accident Reporting*. Glasgow, UK: University of Glasgow Press. [17]

Jones, A. & Wreathall, J. (2000). Leading indicators of human performance – the story so far. *6th Annual Human Performance/Root Cause/Trending Conference*, Philadelphia, PA. [5]

Klaene, B.J. & Sanders, R.E. (2008). *Structural Firefighting: Strategies and Tactics* (2nd ed.). Sudbury, MA: Jones & Bartlett Publishers. [10]

Klein, G.A. & Armstrong, A.A. (2005). Critical decision method. In N. Stanton, A. Hedge, K. Brookhuis, E. Salas & H.W. Hendrick (eds), *Handbook on Human Factors and Ergonomics Methods*. Boca Raton, FL: CRC Press (pp. 35.31–35.38). [3]

Klein, G.A. (1998). *Sources of Power: How People Make Decisions*. Cambridge, MA: MIT Press. [8]

Klein, G.A., Feltovich, P.J., Bradshaw, J. & Woods, D.D. (2005). Common ground and coordination in joint activity. In W. Rouse & K. Boff (eds), *Organizational Simulation*. Chichester, UK: John Wiley & Sons. (pp. 139–84). [4]

Klein, G. A., Pliske, R., Crandall, B. & Woods, D. D. (2005). Problem detection. *Cognition, Technology & Work*, 7, 1, 14–28. [12]

Kline, P. (1994). *An Easy Guide to Factor Analysis*. London, UK: Routledge. [11]

Kluger, M.T., Tham, E.J., Coleman, N.A., Runciman, W.B. & Bullock, M.F. (2000). Inadequate pre-operative evaluation and preparation: a review of 197 reports from the Australian Incident Monitoring Study. *Anaesthesia*, 55(12), 1173–78. [16]

Kontogiannis, T. (1999). Training effective human performance in the managing of stressful emergencies. *Cognition Technology and Work*, 1, 7–24. [8]

Krauss, R.M. & Fussel, S.R. (1990). Mutual Knowledge and Communicative Effectiveness. In J. Galegher et al. (eds), *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work*. Hillsdale: IEA Lawrence Erlbaum Associates (pp. 111–45). [16]

Lagadec, P. (2004). Understanding the French 2003 heat wave experience: Beyond the heat, a multi-layered challenge. *Journal of Contingencies and Crisis Management*, 12(4), 160–9. [16]

Lamond, N. & Dawson, D. (1999). Quantifying the performance impairment associated with fatigue. *Journal of Sleep Research*, 8(4), 255–62. [6]

Lanir, Z. (1986). *Fundamental Surprise*. Eugene, Oregon: Decision Research. [1] [2]

Lave, J. & Wenger, E. (1991). *Situated Learning. Legitimate Peripheral Participation.* Cambridge, MA: Cambridge University Press. [17]

Le Coze, J.-C. (2005). Are organisations too complex to be integrated in technical risk assessment and current safety auditing? *Safety Science*, 43, 613–38. [18]

Leape, L.L., Brennan, T.A., Laird, N., Lawthers, A.G., Localio, A.R., Barnes, B.A., et al. (1991). The nature of adverse events in hospitalized patients. *Results of the Harvard Medical Practice Study II,* 324, (pp. 377–84). [17]

Leplat, J. (1988). Task complexity in work situations. In L.P. Goodstein, H.B. Andersen & S.E. Olsen (eds), *Task, Errors and Mental Models*. London, UK: Taylor and Francis (pp. 105–115). [3]

Leplat, J. (1991). Organization of Activity in Collective tasks. In J. Rasmussen, B. Brehmer, & J. Leplat (eds), *Distributed Decision Making: Cognitive Models for Cooperative Work*. New York: J. Wiley & Sons Ltd (pp. 51–71). [16]

Leveson, N.G. (2002). *Systems Safety Engineering: Back to the Future.* http://sunnyday.mit.edu/book2.pdf, accessed September 2003. [18]

Lewis, D. (1969). *Convention: A Philosophical Study*. Cambridge, MA: Harvard University Press. [16]

Lumpé, M.-P. (2008). *Leadership and Organization in the Aviation Industry*. Aldershot, UK: Ashgate. [18]

Lundberg, J., Rollenhagen, C. & Hollnagel, E. (2009). What-you-look-for-is-what-you-find – The consequences of underlying accident models in eight accident investigation manuals. *Safety Science*, 47(10), 1297–1311. [14] [18] [Epilogue]

Malakis, S., Kontogiannis, T. & Kirwan, B. (2010a). Managing emergencies and abnormal situations in air traffic control (part I): Taskwork strategies. *Applied Ergonomics*, doi:10.1016/j.apergo.2009.12.019. Published online. [8]

Malakis, S., Kontogiannis, T. & Kirwan, B. (2010b). Managing emergencies and abnormal situations in air traffic control (part II): Teamwork strategies. *Applied Ergonomics*, doi:10.1016/j.apergo.2009.12.018. Published online. [8]

Manganelli, S. & Engle, R.F. (2001). *Value at Risk in Finance*. Working Paper No 75, European Central Bank (ECB) Working Paper Series, August. [13]

Marais, K., Dulac, L. & Leveson, N.G. (2004). *Beyond Normal Accidents and High Reliability Organizations: The Need for an Alternative Approach to Safety in Complex Systems,* in proceedings of the MIT ESD Symposium (online), retrieved January 2007 from http://sunnyday.mit.edu/ papers/hro.pdf. [18]

Mataric, M.J. (1992). Minimizing complexity in controlling a mobile robot population. In *Proceedings of the 1992 International Conference on Robotics and Automation*, Los Alamitos: IEEE Computer Society Press. [16]

McDonald, N. (2006). Organizational Resilience and Industrial Risk. In E. Hollnagel, D.D. Woods. and N. Leveson (eds), *Resilience Engineering Concepts and Precepts*. Farnham, UK: Ashgate (pp. 155–80). [12]

Mendonça, D. (2008). Measures of Resilient Performance. In C. Nemeth, C.E. Hollnagel & S.W.A. Dekker, (eds), *Resilience Engineering Perspectives Vol.1: Remaining Sensitive to the Possibility of Failure*. Aldershot, UK: Ashgate (pp. 29–38). [11]

Merton, R.C. (1995). A functional perspective of financial intermediation. *Financial Management*, 24(2), 21–41. [13]

Merton, R.C. & Bodie, Z. (1995). A conceptual framework for analysing the financial environment. In D.B. Crane, K.A. Froot, S.P. Mason, A.F. Perold & R.C. Merton (eds), *The Global Financial System*. Cambridge, MA: Harvard Business School Press (pp. 3–31). [13]

Merton, R.C. & Bodie, Z. (2005). Design of financial systems: Towards a synthesis of function and structure. *Journal of Investment Management*, 3(1), 1–23. [13]

Miller, A. & Xiao, Y. (2007). Multi-level strategies to achieve resilience for an organisation operating at capacity: a case study at a trauma centre. *Cognition, Technology & Work*, 9, 51–66. [3]

Ministerie van Verkeer en Waterstaat (Ministry of Transportation), (2008). *Ontwikkeling van Schiphol en haar omgeving op middellange termijn.* Alderstafel, Ministerie van Verkeer en Waterstaat. [15]

Mintzberg, H. (1996). *The Rise and Fall of Strategic Planning.* New York, USA: Free Press. [1] [2]

Mollo, V. & Falzon, P. (2004). Auto- and allo-confrontation as tools for reflective activities. *Applied Ergonomics*, 35(6), 531–40. [3]

Morel, G., Amalberti, R. & Chauvin, C. (2008). Articulating the differences between safety and resilience: The decision-making process of professional sea-fishing skippers. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 50(1), 1–16. [12]

NASB, (1994). *Aircraft Accident report 92-11. El-Al Flight 1862*. Netherlands Aviation Safety Board. [15]

National Transportation Safety Board (1994). *Uncontrolled Collision with Terrain. American International Airways Flight 808*. NTSB/AAR-94/04, National Transportation Board, Washington DC. [6]

National Transportation Safety Board (NTSB) (1999). *Evaluation of U.S. Department of Transportation Efforts in the 1990s to Address Operator Fatigue*. NTSB/SR-99/01. National Transportation Safety Board, Washington D.C. [6]

Neri, D.F. & Nunneley, S.A. (2004). Proceedings of the fatigue and performance modeling workshop, June 13–14 2002, Seattle, WA. *Aviation, Space and Environmental Medicine*, 75(3, Section II), A1–A199. [6]

Norman, D.A. (1987). *The Psychology of Everyday Things*. New York: Basic Books. [16]

Norman, D.A. (1990). The 'problem' of automation: Inappropriate feedback and interaction, not 'over-automation.' *Philosophical Transactions of the Royal Society of London, B*, 327, 585–93. [10]

Nyssen, A.S. (2008). Coordination in hospitals: organized or emergent process? Towards the idea of resilience as the agents,' groups,' systems' capacity to project themselves into future. In E. Hollnagel, F. Pieri & E. Rigaud (eds), *3rd International Symposium on Resilience Engineering*. October 28–30, France: Antibes – Juan-les-Pins. [3]

Nyssen, A.S. & Javaux, D. (1996). Analysis of synchronization constraints and associated errors in collective work environments. *Ergonomics*, 39, 1249–64. [16]

Ombredane, A. & Faverge, J.-M. (1955). *L'Analyse du Travail*. Paris: PUF. [3]

Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. New York: Cambridge University Press, 1990. [10]

Ostrom, E. (1999). Coping with tragedies of the commons. *Annual Review of Political Science*, 2, 493–535. [10]

Pariès, J. (1996). Evolution of the aviation safety paradigm: Towards systemic causality and proactive actions. In B. Hayward & A. Lowe (eds.) *Proceedings of the Australian Aviation Psychology Symposium 1995: Applied Aviation Psychology: Achievement, Change and Challenge*. Manly, Australia (pp. 39–49). [18]

Pariès, J. (1999). *Shift in Aviation Safety Paradigm is Key to Future Success in Reducing Air Accidents.* Presentation to the ICAO Regional Symposium on Human Factors and Aviation Safety, Santiago du Chili. *ICAO Journal*, 54(5). Montréal, Canada. [18]

Pariès, J. & Amalberti, R. (2000). Aviation safety paradigms and training implications. In N.S. Sarter & R. Amalberti (eds), *Cognitive Engineering in the Aviation Domain*. New Jersey: Lawrence Erbaum Associates (pp. 253–86). [1] [8]

Parliamentary Inquiry, (1999). *Een beladen vlucht.* Enquetecommissie Vliegramp Bijlmermeer. [15]

Patterson, E.S. & Woods, D.D. (2001). Shift changes, updates, and the on-call model in space shuttle mission control. Computer supported cooperative work. *The Journal of Collaborative Computing*, 10, 3–4, 317–46. [12]

Pavard, B. (1994). *Système Coopératifs: de la Modélisation à la Coopération*. Toulouse: Octares. [16]

Perrenoud, P. (1999). Gestion de l'imprévu, analyse de l'action et construction de compétences. *Education Permanente*, 140(3), 123–44. [3]

Perrow, C. (1984). *Normal Accidents: Living with High-Risk Technologies*. New York, USA: Basic books. [10] [12] [18]

Pettersen, K.A. (2006). Operational problem solving in aviation - the role of social and organisational factors in safety. In G. Soares & Sio (eds), *Safety and Reliability for Managing Risk.* London, Taylor & Francis Group. [12]

Pettersen, K.A. (2008). The Social Production of Safety. Theorising the Human Role in Aircraft Line Maintenance. *PhD thesis, University of Stavanger*, No.59, December 2008 (pp. 407–412). [12]

Pettersen, K.A. & Aase, K. (2008). Explaining safe work practices in aviation line maintenance. *Safety Science*, 42, 10–19. [12]

Piaget, J. (1967). *Biologie et Connaissance.* Paris: Gallimard. [1]

Piaget, J. (1967/1992). *Biologie et Connaissance*. Lausanne: Delachaux et Niestlé. (Première édition: 1967, Paris: Gallimard). [16]

Potter, J. & Wetherell, M. (1987). *Discourse and Social Psychology: Beyond Attitudes and Behaviour*. Thousand Oaks, CA: Sage Publications. [18]

PRC, (2006). *Evaluatie Functioneren VACS*. Eindrapport. Policy Research Corporation, Antwerp, Belgium. [15]

RAND Europe, (1993). Airport Growth and Safety. A Study of the External Risks of Schiphol Airport and Possible Safety-Enhancement measures. EAC Rand, Delft, the Netherlands. [15]

Rasmussen, J. (1997). Risk management in a dynamic society: a modelling problem. *Safety Science*, 27(2–3), 183–213. [3] [12] [13]

Reason, J.T. (1990). *Human Error*. Cambridge, UK: Cambridge University Press. [12] [17]

Reason, J.T. (1997). *Managing the Risks of Organizational Accidents*. Aldershot, UK: Ashgate. [17]

Reason, J.T. (1998). *Managing the Risks of Organizational Accidents*. Brookfield, VT: Ashgate. [5]

Reason, J.T. (2009). *The Human Contribution*. Aldershot, UK: Ashgate. [7]

Reason, J.T. & Hobbs, A. (2003). *Managing Maintenance Error*. Aldershot, UK: Ashgate. [7]

Reynolds, C.W. (1987). Flocks, herds, and schools: A distributed behavioral models. *Computer Graphics*, 21, 25–34. [16]

Rochlin, G.I. (1999). Safe operation as a social construct. *Ergonomics*, 42(11), 1549–60. [10]

Roelen, A. (2008). Causal Risk Models of Air Transport. Comparison of user Needs and Model capabilities. *Doctoral Thesis. Delft University of Technology*. [15]

Rosenthal, U. (1999). *International Conference on the Future of European Crisis Management.* The Hague, November 7–9. Challenges of crisis management in Europe. [15]

Rudolph, J.W., Morrison, J.B. & Carroll, J.S. (2009). The dynamics of action-oriented problem solving: Linking interpretation and choice. *Academy of Management Review*, 34(4), 733–56. [10]

Salas, E., Bowers, C.A. & Rhodenizer, L. (1998). It is not how much you have but how you use it: Toward a rational use of simulation to support aviation training. *International Journal of Aviation Psychology*, 8, 97–208. [8]

Salas, E., Sims, D.E. & Burke, C.S. (2005). Is there a 'big five' in teamwork? *Small Group Research*, 36, 555–99. [8]

Savoyant, A. & Leplat, J. (1983). Statut et fonction des communications dans l'activité des équipes de travail (Statut and function of the communications in the activities of the workteams). *Psychol. fr*., 28(3), 247–53. [16]

Scheffer, M., Bascompte, J., Brock, W.A., Brovkin, V., Carpenter, S.R., Dakos, V., Held, H., van Nes, E.H., Rietkerk, M. & Sugihara, G. (2009). Early-warning signals for critical transitions. *Nature*, 461(7260), 53–59. [10]

Schmidt, R.H. & Tyrell, M (2004). What constitutes a financial system in general and the German financial system in particular? In J.P. Krahnen & R.H. Schmidt (eds), *The German Financial System*. Oxford, UK: Oxford University Press (pp. 19–67). [13]

Schwarz, N. (2007). Attitude construction: Evaluation in context. *Social Cognition*, 25, 638–56. [18]

Shattuck, L.G. & Woods, D.D. (1997). Communication of intent in distributed supervisory control system. In *Proceedings of the 41st Annual Meeting of the Human Factors and Ergonomics Society*, September 1997. [12]

Signal, T.L., Ratieta, D. & Gander, P.H. (2008). Flight crew fatigue management in a more flexible regulatory environment: an overview of the New Zealand aviation industry. *Chronobiology International*, 25(2–3), 373–88. [6]

Simon, H. (1982). *Models of Bounded Rationality: Behavioral Economics and Business Organization (*Vols 1 & 2). Cambridge, MA: The MIT Press. [2]

Simpson, R. (1996). Neither clear nor present: The social construction of safety and danger. *Sociological Forum*, 11(3), 549–62. [18]

Sklet, S. (2002*). Methods for Accident Investigation*. Department of Production and Quality Engineering. NTNU Norwegian University of Science and technology. [15]

Snook, S.A. (2000). *Friendly Fire.* New Jersey: Princeton University Press. [12]

Spencer, M.B., Robertson, K.A., Cabon, P., Mollard, R., Åkerstedt, T., Guillberg, M., Simon, R., Valk, P., Samel, A. & Gundel, A. (2002). *Modelling of Aircrew Alertness in Future Ultra Long-Range Schedules, Based on a City Pair*. QinetiQ Report No QINETIQ/CHS/P&D/CR020047/1.1, February 2002. [6]

Spencer, M.B. & Robertson, K.A. (2007). The application of an alertness model to ultra-long-range civil air operations. *Somnologie*, 11, 159–66. [6]

Stanton, N.A., Salmon, P., Walker, G.H., Baber, C. & Jenkins, D.P. (2005). *Human Factors Methods – A Practical Guide for Engineering and Design*. Aldershot, UK: Ashgate. [3]

Steele, K.R. & Pariès, J. (2007). Barriers to safety innovation: Experiences applying the 'Safety Model Based Analysis' approach in European aviation. *Proceedings of the 14th International Symposium on Aviation Psychology*. Ohio, USA. [18]

Steenhuisen, B. & Van Eeten, M. (2008). Invisible trade-offs of public values: Inside Dutch railways. *Public Money & Management*, June 2008, 147–52. [15]

Stewart, S. (2007). An integrated system for managing fatigue risk within a low cost carrier. In *Proceedings of the International Aviation Safety Seminar*, Flight Safety Foundation, October 23–26, Paris, France. [6]

Stoop, J. (1990). Safety and the Design Process. *Doctoral Thesis*. *Delft University of Technology*, April 1990. [15]

Stoop, J. (1997). Airport growth and safety: improvement of the external and internal risks of airports. *International Aviation Safety Conference, IASC-79, August 27–29.* The Netherlands: Rotterdam Airport. [15]

Stoop, J. (2003). Critical size events: a new tool for crisis management resource allocation? *Safety Science*, 41, 465–80. [15]

Stoop, J. (2009). Before, during and after the event; the Boeing 747 case study. 36th ESReDA *Lessons Learned from Accident Investigations*, 2–3 June 2009, Seminar Coimbra, Portugal. [15]

Stoop, J. & Dekker, S.W.A. (2007). Are safety investigates proactive? *33rd ESReDA Seminar Future Challenges of Accident Investigation.* Ispra, Italy, November 13–14. [15]

Stoop, J., Baggen, J.H., De Kroes, J.L., Vleugel, J.M. & Vrancken, J.L.M. (2007). HSL safety signalling system ERTMS. An independent investigation into the usefulness of adapting the ERTMS safety signalling system. *Commissioned by the Research and Verification Department of the Dutch Parliament*. Delft University of Technology, 23 May 2007 (In Dutch with English summary). [15]

Strohschneider, S. & Gerdes, J. (2004). MS: ANTWERPEN: Emergency management training for low-risk environments. *Simulation & Gaming, 35*. [4]

Sundström, G.A. & Hollnagel, E. (2006). Learning how to create resilience in business systems. In E. Hollnagel, D.D. Woods & N. Leveson (eds), *Resilience Engineering. Concepts and Precepts*. Aldershot, UK: Ashgate. [13]

Tabachnick, B. & Fidell, L. (2007). *Using Multivariate Statistics (5th Edition)*. Boston, USA: Allyn and Bacon. [11]

Thomas, M. J. W., Petrilli, R.M., Lamond, N., Dawson, D. & Roach, G.D. (2006). Australian Long Haul Fatigue Study. In *Enhancing Safety Worldwide: Proceedings of the 59th Annual IASS*. Alexandria, Virginia, U.S.: Flight Safety Foundation. [6]

Tjørhom, B. & Aase, K. (2007). Safety and changes in the Norwegian aviation transport system – What is the role of the legislator and the regulator? In T. Aven and J.E. Vinnem. (eds), *Risk Reliability and Societal Safety. Volume 3*, 2143–49. Taylor & Francis. [12]

Tjørhom, B. & Aase, K. (2010). The role of complexity in accident investigation practice. *International Journal of Emergency Management*, 7(2), 167–189. [12]

Touchman, B.W. (1985). *The March of Folly: From Troy to Vietnam*. New York: Ballantine Books. [14]

Turner, B. (1978). *Man-Made Disasters*. London: Wykeham. [18]

Turoff, M. & Linstone, H.A. (eds), (1975). *The Delphi Method – Techniques and Applications*. (www.is.njit.edu/pubs/delphibook/) London, UK: Addison-Wesley. [11]

Uhr, C. (2009). Multi-organizational Emergency Response Management – A Framework for Further Development. Department of Fire Safety Engineering and Systems Safety. *Lund, Lund University. Doctoral thesis*. [4]

Van Binsbergen, A., Van Eeghen, M., Polinder, B., Stoop, J., Wiegmans, B. & Zigterman, L. (2008). *Quick-scan Double Rail*. TRAIL Research School, Delft University of Technology, June 2008. [15]

Van Dam, S., Mulder, M. & Van Paassen, M.M. (2008). Ecological interface design of a tactical airborne separation assistance tool. *Journal of IEEE Transactions on Systems, Man, and Cybernetics*, 38(6), 1221–33. [15]

Van der Shaaf, T.W., Lucas, D.A. & Hale, A.R. (1991). *Near Miss Reporting as a Safety Tool*. Oxford, UK: Butterworth-Heinemann. [17]

Van Dongen, H.A., Maislin, G., Mullington, J. & Dinges, D.F. (2003). The cumulative cost of additional wakefulness: dose-response effects on neurobehavioral functions and sleep physiology from chronic sleep restriction and total sleep deprivation. *Sleep*, 26, 117–26. [6]

Van Eijndhoven, J. (2008). Increasing flexibility by combining business processes with business rules. *M.Sc. Thesis Business Information Technology*. Enschede, The Netherlands: TNO-ICT. [15]

Vaughan, D (2006). The social shaping of commission reports. *Sociological Forum*, 21, 2, 291–306. [12]

Vaughan, D. (1996). *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. Chicago: University of Chicago Press. [12]

Vincent, C.A. (2006). *Patient Safety*. Amsterdam, The Netherlands: Churchill Livingstone (Elsevier). [17]

Von Bertalanffy, L. (1975). *Perspectives on General Systems Theory*. New York: George Braziller. [13]

Wears, R.L. & Woods, D.D. (2007). Always adapting. *Annals of Emergency Medicine*, 50(5), 517–19. [10]

Wears, R.L., Perry, S. & McFauls, A. (2006). 'Free fall' – A case study of resilience, its degradation, and recovery in an emergency department. In E. Hollnagel & E. Rigaud (eds), *2nd International Symposium on Resilience Engineering*. 8–10 November, France, Juan-les-Pins. [3]

Wears, R.L., Perry, S., Anders, S. & Woods, D.. (2008). Resilience in the emergency department. In E. Hollnagel, C. Nemeth & S.W.A. Dekker (eds), *Resilience Engineering Perspectives: Remaining Sensitive to the Possibility of Failure* (Vol. 1, pp. 193–210). Aldershot, UK: Ashgate. [3]

Weick, K. (1993). Collapse of sensemaking in organizations: The Mann Gulch Disaster. *Administrative Science Quarterly*, 38, 628–52. [2] [3]

Weick, K. & Sutcliffe, K. (2001). *Managing the Unexpected*. San-Francisco: John Wiley & Sons, Inc. [16]

Westrum, R. (1993). Cultures with requisite imagination In J.A. Wise, V.D. Hopkin & P. Stager (eds), *Verification and Validation of Complex Systems: Human Factor Issues*. New York, NY: Springer-Verlag. [17]

Westrum, R. (1999). Faint hearts and faint signals – how organizations manage signs of trouble. *1999 Workshop of the Center for Human Performance in Complex Systems*, Madison, WI, University of Wisconsin. [5]

Westrum, R. (2006). A typology of resilience situations. In E. Hollnagel, D.D. Woods & N.G. Leveson (eds), *Resilience Engineering: Concepts and Precepts*. Aldershot, UK: Ashgate (pp. 55–65). [Prologue] [1] [2] [3] [13]

Wilson, K.A., Salas, E., Priest, H.A. & Andrews, D. (2007). Errors in the heat of battle: Taking a closer look at shared cognition breakdowns through teamwork. *Human Factors*, 49, 243–56. [8]

Woods, D.D. (2004). *Creating Foresight: Lessons for enhancing resilience from Columbia*. Retrieved January 17th, 2006 from http://csel.eng.ohio-state.edu/woods/space/Create%20foresight%20Col-draft.pdf. [18]

Woods, D.D. (2005). Creating foresight: Lessons for resilience from Columbia. In M. Farjoun and W.H. Starbuck (eds), *Organization at the Limit: NASA and the Columbia Disaster*. Hoboken, NJ: Blackwell (pp. 289–308).[10] [12]

Woods, D.D. (2006a). Essential characteristics of resilience. In E. Hollnagel, D.D. Woods, N.G. Leveson (eds), *Resilience Engineering: Concepts and Precepts*. Aldershot, UK: Ashgate (pp. 21–33). [1] [2] [3] [10] [11] [12] [17]

Woods, D.D. (2006b). How to design a safety organization: Test case for resilience engineering. In E. Hollnagel, D.D. Woods and N. Leveson (eds), *Resilience Engineering Concepts and Precepts*. Farnham, UK: Ashgate (pp. 315–24). [12]

Woods, D.D. (2009a). Escaping failures of foresight. *Safety Science*, 47(4), 498–501. [10]

Woods, D.D. (2009b). Fundamentals to engineer resilient systems: How human adaptive systems fail and the quest for polycentric control architectures. Keynote presentation, *2nd International Symposium on Resilient Control Systems,* Idaho Falls, ID, August 11–13 2009 (https://secure.inl.gov/isrcs2009/default.aspx accessed September 8, 2009). [10]

Woods, D.D. & Cook, R.I. (2002). Nine steps to move forward from error. *Cognition, Technology, and Work*, 4(2), 137–44. [14]

Woods, D.D. & Cook, R.I. (2006). Incidents – markers of resilience or brittleness? In E. Hollnagel, D.D. Woods & N. Leveson (eds), *Resilience Engineering: Concepts and Precepts*. Aldershot, UK: Ashgate (pp. 69–76). [3] [10]

Woods, D.D., Dekker, S.W.A., Cook, R.I., Johannesen, L.L. & Sarter, N.B. (2010). *Behind Human Error* (2nd Edition). Farnham, UK: Ashgate. [10]

Woods, D.D. & Hollnagel, E. (2006). *Joint Cognitive Systems: Patterns in Cognitive Systems Engineering.* Boca Raton, FL: CRC Press. Taylor & Francis Group. [8] [10]

Woods, D.D., Johannesen, L.J., Cook, R.I. & Sarter, N.B. (1994). *Behind Human Error: Cognitive Systems, Computers, and Hindsight.* Wright-Patterson Air Force Base, OH: Crew Systems Ergonomics Information Analysis Center. [8] [12]

Woods, D.D. & Patterson, E.S. (2000). How unexpected events produce an escalation of cognitive and coordinative demands. In P.A. Hancock and P. Desmond (eds), *Stress, Workload and Fatigue*. Hillsdale NJ: Lawrence Erlbaum (pp. 290–302). [4]

Woods, D.D., Patterson, E.S. & Cook, R.I. (2007). Behind human error: Taming complexity to improve patient safety. In P. Carayon. (ed.), *Handbook of Human Factors and Ergonomics in Health Care and Patient Safety.* Lawrence Erlbaum Associates. New Jersey: Mahwah. [8]

Woods, D.D., Patterson, E.S. & Roth, E.M. (2002). Can we ever escape from data overload? A cognitive systems diagnosis. *Cognition, Technology & Work*, 14, 22–36. [4]

Woods, D.D. & Sarter, N.B. (2000). Learning from automation surprises and 'going sour' accidents. In N. Sarter & R. Amalberti (eds), *Cognitive Engineering in the Aviation Domain*. Hillsdale NJ: Erlbaum (pp. 327–54). [10 [14]]

Woods, D.D. & Shattuck, L.G. (2000). Distant supervision – Local action given the potential for surprise. *Cognition, Technology & Work*, 2, 242–45. [10] [12]

Woods, D.D. & Wreathall, J. (2008). Stress-strain plots as a basis for assessing system resilience. In E. Hollnagel, C. Nemeth & S.W.A. Dekker (Eds.) *Remaining Sensitive to the Possibility of Failure*. Aldershot, UK: Ashgate. (pp. 143–158). [7] [3] [10]

Wreathall, J. (2006). Properties of resilient organizations: An initial view. In Hollnagel, E., Woods, D. D., Leveson, N. (eds), *Resilience Engineering: Concepts and Precepts*. Aldershot, UK: Ashgate. (pp. 275–85) [11] [5]

Wreathall, J. (2009). Leading? Lagging? Whatever! *Safety Science*, 47(4), 493–94. [5]

Wreathall, J. & Jones, A. (2000). Leading indicators of human performance – the story so sar. *6th Annual Human Performance/Root Cause/Trending Conference*, Philadelphia, PA. [5]

Wreathall, J. & Merritt, A.C. (2003). Managing Human Performance in the Modern World: Developments in the US Nuclear Industry. In Edkins, G. & Pfister, P. (eds), *Innovation and Consolidation in Aviation*. Aldershot, UK: Ashgate. (pp. 159–70). [7]

Zaller, J. & Feldman, S. (1992). A simple theory of survey response: Answering questions versus revealing preferences. *American Journal of Political Science*, 36, 579–616. [18]

Zhou, T., Carlson, J.M. & Doyle, J. (2005). Evolutionary dynamics and highly optimized tolerance. *Journal of Theoretical Biology*, 236, 438–47. [10]

Zimmermann, K. (2009). Unpublished interview data. [18]

This page has been left blank intentionally

Author Index

This page has been left blank intentionally

Subject Index