

**The Organisational Precursors to  
Human Automation Interaction  
Issues in Safety-Critical Domains**

**The Case of an Automated Alarm System from the Air  
Traffic Management Domain**

By

Simone Rozzi

Dissertation submitted in partial fulfilment  
of the requirements for the degree of  
Doctor of Philosophy

September 2015

**School of Science and Technology,  
Middlesex University, London, UK**

© Simone Rozzi 2015

# Abstract

Much has been written about the side effects of automation in complex safety-critical domains, such as air traffic management, aviation, nuclear power generation, and healthcare. Here, human factors and safety researchers have long acknowledged that the potential of automation to increase cost-effectiveness, quality of service and safety, is accompanied by undesired side effects or issues in human automation interaction (HAI). Such HAI issues may introduce the potential for increased confusion, uncertainty, and frustration amongst sharp end operators, i.e. the users of automation. These conditions may result in operators to refuse to use the automation, in impaired ability of operators to control the hazardous processes for which they are responsible, and in new, unintended paths to safety failure.

The present thesis develops a qualitative framework of the *organisational precursors to HAI issues* (OPHAI) that can be found in safety-critical domains. Organisational precursors denote those organisational and managerial conditions that, although distant in time and space from the operational environment, may actually influence the quality of HAI found there. Such precursors have been extensively investigated by organisational safety (OS) scholars in relation to the occurrence of accidents and disasters—although not HAI issues. Thus, the framework's development is motivated by the intent to explore the theoretical gap lying at the intersection between the OS area and the current perspectives on the problem—the human computer interaction (HCI) and the system lifecycle ones. While considering HAI issues as a design problem or a failure in human factors integration and/or safety assurance respectively, both perspectives, in fact, ignore, the organisational roots of the problem.

The OPHAI framework was incrementally developed based on three qualitative studies: two successive, historical, case studies coupled with a third corroboratory expert study. The first two studies explored the organisational precursors to a known HAI issue: the nuisance alert problem relative to an automated alarm system from the air traffic management domain. In particular, the first case study investigated retrospectively the organisational response to the nuisance alert problem in the context of the alarm's implementation and improvement in the US between 1977 and 2006. The second case study has a more contemporary focus, and examined at the organisational response to the same problem within two European Air Navigation Service Providers between 1990

and 2010. The first two studies produced a preliminary version of the framework. The third study corroborated and refined this version by subjecting it to the criticism from a panel of 11 subject matter experts.

The resulting framework identifies three classes of organisational precursors: (i) the organisational assumptions driving automation adoption and improvement; (2) the availability of specific organisational capabilities for handling HAI issues; and (3) the control of implementation quality at the boundary between the service provider and the software manufacturer. These precursors advance current understanding of the organisational factors involved in the (successful and problematic) handling of HAI issues within safety-critical service provider organisations. Its dimensions support the view that HAI issues can be seen as an organisational phenomenon—an organisational problem that can be the target of analysis and improvements complementary to those identified by the HCI and the system lifecycle perspectives.

# Acknowledgments

Academic supervisors:

Dr Bob Fields  
Dr John Watt  
Prof Darren Dalcher  
Dr Paola Amaldi

EUROCONTROL staff:

Dr Barry Kirwan  
Andy Kilner  
Ben Bakker  
Stan Drozdowsky  
Hans Wagemans

Sponsor:

EUROCONTROL Experimental Centre<sup>1</sup>,  
Brétigny-sur-Orge.

Colleagues that generously  
helped out with the review of  
the manuscript:

Dr Luca Save  
Ronish Joyekurn  
Dr Marco Mazzeo

Research participants:

All the research participants from  
the SPIN network and  
the organisations  
accessed in Study 2

Special thanks to:

Lorenzo Bocedi  
Prof Richard Comley  
Prof Martin Loomes  
Prof William Wong  
Prof Giampietro Gobo  
Prof Vu Duong

Fellow doctoral students and  
colleagues at EUROCONTROL  
Experimental Centre, Paris,  
France:

Stefano Tiberia  
Marinella Leone  
Luca Bellesia  
Dr Andrea Ranieri  
Maurizio Romano  
Dr Magnus Axholt  
Dr Stephen Peterson  
Dr Ella Pinska

Important people this work is  
passionately dedicated to:

My beautiful, brilliant and  
outrageously loving and  
supportive wife, Valeria  
my son, Leonardo  
my family

---

<sup>1</sup> Founded and hosted the first three years of the PhD.

# Table of Contents

<b>Abstract</b> .....	<b>iii</b>
<b>Acknowledgments</b> .....	<b>v</b>
<b>Table of Contents</b> .....	<b>vi</b>
<b>List of Abbreviations</b> .....	<b>x</b>
<b>List of Tables</b> .....	<b>xii</b>
<b>List of Figures</b> .....	<b>xiv</b>
<b>CHAPTER 1. INTRODUCTION</b> .....	<b>1</b>
1.1. RESEARCH GAP .....	3
1.2. AIM AND OBJECTIVES .....	5
1.3. THESIS STRUCTURE .....	5
<b>CHAPTER 2. HAI ISSUES: CURRENT THEORETICAL PERSPECTIVES</b> .....	<b>7</b>
2.1. CHAPTER INTRODUCTION .....	7
2.2. THE HUMAN COMPUTER INTERACTION PERSPECTIVE .....	8
2.2.1. Human information processing .....	8
2.2.2. Distributed cognition .....	9
2.2.3. Activity theory .....	10
2.2.4. Computer supported collaborative work .....	11
2.2.5. Cognitive system engineering .....	12
2.2.6. Critique of the HCI perspective .....	14
2.3. SYSTEM LIFECYCLE PERSPECTIVE .....	15
2.3.1. User centred design .....	15
2.3.2. Human factors integration .....	16
2.3.3. System safety .....	16
2.3.4. Critique of the system lifecycle perspective .....	17
2.4. CHAPTER CONCLUSIONS .....	19
<b>CHAPTER 3. THE ORGANISATIONAL SAFETY PERSPECTIVE</b> .....	<b>20</b>
3.1. CHAPTER INTRODUCTION .....	20
3.2. THE ORGANISATIONAL SAFETY PERSPECTIVE AND ORGANISATIONAL DRIFT INTO FAILURE .....	21
3.2.1. A by-product of normal organisational and administrative activity .....	22
3.2.2. Induced by decisions located at different levels of society at different points in time .....	25
3.2.3. An incremental process .....	28
3.2.4. An interpretive phenomenon .....	32
3.2.5. Induced by the intrinsic complexity of complex, safety-critical systems .....	33
3.3. AVOIDING ORGANISATIONAL DRIFT INTO FAILURE .....	35
3.4. DISCUSSION: OUTLINING THE RESEARCH GAP AND OBJECTIVE .....	37
3.5. CHAPTER CONCLUSIONS .....	40

<b>CHAPTER 4. RESEARCH STRATEGY .....</b>	<b>41</b>
4.1. CHAPTER INTRODUCTION.....	41
4.2. OVERVIEW .....	42
4.3. RATIONALE .....	44
4.3.1. Rationale for grounding the research on the case study approach .....	44
4.3.2. Boundaries of the research.....	46
4.3.3. Selected application case: the MSAW .....	47
4.3.4. Validation and generalisability .....	48
4.4. DATA AND DATA COLLECTION .....	50
4.4.1. Study 1 .....	51
4.4.1.1. Data sources.....	51
4.4.1.2. Data bases accessed.....	52
4.4.2. Study 2 .....	53
4.4.2.1. Sample of organisations .....	53
4.4.2.2. Organisational access .....	54
4.4.2.3. Data collection .....	55
4.4.2.4. Data sources.....	56
4.4.3. Study 3 .....	57
4.4.3.1. Recruiting of experts.....	57
4.4.3.2. Data Collection .....	57
4.4.3.3. Questionnaire .....	58
4.4.3.4. Profile of the SME group .....	59
4.5. DATA ANALYSIS .....	62
4.5.1. Study 1 .....	64
4.5.1.1. First-level coding for Study 1 .....	65
4.5.1.2. Second-level coding for Study 1 .....	68
4.5.2. Study 2 .....	69
4.5.2.1. First-level coding for Study 2.....	69
4.5.2.2. Second-level coding for Study 2 .....	70
4.5.3. Study 3 .....	74
4.6. WITHIN-STUDY VALIDATION STRATEGIES .....	75
4.7. CHAPTER CONCLUSIONS .....	77
<b>CHAPTER 5. STUDY 1 RESULTS .....</b>	<b>78</b>
5.1. CHAPTER INTRODUCTION.....	78
5.2. BACKGROUND .....	79
5.2.1. US MSAW .....	79
5.2.2. The organisational context of analysis .....	79
5.2.2.1. FAA Response to NTSB Safety Recommendations .....	80
5.2.2.2. Ensuing correspondence exchange.....	80
5.3. FIRST-LEVEL CODING RESULTS .....	82
5.3.1. Coding framework 1.1: MSAW-related concerns identified by NTSB.....	82
5.3.2. Coding framework 1.2: Areas of changes requested by the NTSB to address the “controller lack-of-response” concern.....	85
5.3.3. Coding Framework 1.3 and 1.4: analysis of the contrasted changes .....	88
5.3.3.1. RC1: Disambiguate between the MSAW and CA aural alerts .....	88

5.3.3.2.	RC7: Amend existing regulations to make mandatory the transmission of MSAW alert to the pilot.....	92
<b>5.4.</b>	<b>SECOND-LEVEL CODING RESULTS: ORGANISATIONAL PRECURSORS IDENTIFICATION.....</b>	<b>97</b>
5.4.1.	History of the nuisance alert problem in the studied context.....	97
5.4.2.	Organisational assumptions about the role of the alarm.....	100
5.4.2.1.	NTSB: viewing the MSAW as a “hazard resolver”.....	100
5.4.2.2.	FAA: viewing the alarm as an “attention director”.....	101
<b>5.5.</b>	<b>CHAPTER CONCLUSIONS.....</b>	<b>102</b>
<b>CHAPTER 6.</b>	<b>STUDY 2 RESULTS.....</b>	<b>103</b>
<b>6.1.</b>	<b>CHAPTER INTRODUCTION.....</b>	<b>103</b>
<b>6.2.</b>	<b>BACKGROUND.....</b>	<b>104</b>
<b>6.3.</b>	<b>FIRST-LEVEL CODING RESULTS: ORGANISATION SPECIFIC FINDINGS.....</b>	<b>105</b>
6.3.1.	Alphasky practices.....	105
6.3.2.	Deltasky practices.....	110
<b>6.4.</b>	<b>SECOND-LEVEL CODING RESULTS: ORGANISATIONAL PRECURSORS IDENTIFICATION.....</b>	<b>114</b>
6.4.1.	OP1. Organisational assumptions driving implementation and improvement.....	116
6.4.2.	OP2. Organisational capability for handling HAI issues.....	117
6.4.3.	OP3. Control over implementation quality at the boundary between the service provider and the software manufacturer.....	118
<b>6.5.</b>	<b>CHAPTER CONCLUSIONS.....</b>	<b>120</b>
<b>CHAPTER 7.</b>	<b>STUDY 3 RESULTS.....</b>	<b>121</b>
<b>7.1.</b>	<b>CHAPTER INTRODUCTION.....</b>	<b>121</b>
<b>7.2.</b>	<b>RESULTS.....</b>	<b>122</b>
7.2.1.	SME feedback on OP1: Organisational Assumptions Driving Automation Implementation And Improvement.....	123
7.2.1.1.	Implications for OP1.....	125
7.2.2.	SME feedback on OP2: Organisational capability for handling HAI.....	126
7.2.2.1.	Implications for OP2.....	130
7.2.3.	SME feedback on OP3: Control over implementation quality at the boundary between the service provider and the software manufacturer.....	132
7.2.3.1.	Implications for OP3.....	134
<b>7.3.</b>	<b>SUMMARY AND DISCUSSION.....</b>	<b>135</b>
<b>7.4.</b>	<b>CHAPTER CONCLUSIONS.....</b>	<b>137</b>
<b>CHAPTER 8.</b>	<b>DISCUSSION AND CONCLUSIONS.....</b>	<b>138</b>
<b>8.1.</b>	<b>CHAPTER INTRODUCTION.....</b>	<b>138</b>
<b>8.2.</b>	<b>THEORETICAL CONSIDERATIONS.....</b>	<b>141</b>
8.2.1.	Comparisons with the HCI perspective.....	141
8.2.2.	Comparison with system lifecycle perspectives.....	142
8.2.3.	Comparison with the OS perspective.....	143
<b>8.3.</b>	<b>METHODOLOGICAL CONSIDERATIONS.....</b>	<b>144</b>
8.3.1.	Considerations on the data used in Study 1.....	144
8.3.2.	Considerations on the value of the organisational trajectory of an HAI issue.....	146
<b>8.4.</b>	<b>PRAGMATIC CONSIDERATIONS.....</b>	<b>147</b>



8.4.1.	Comparison with frameworks for handling HAI issues .....	149
8.4.2.	Comparison with comparable safety investigation methods.....	150
8.4.3.	Comparison with socio-technical frameworks of healthcare IT implementation.....	151
<b>8.5.</b>	<b>GENERALISABILITY .....</b>	<b>152</b>
<b>8.6.</b>	<b>LIMITATIONS &amp; FUTURE WORK.....</b>	<b>154</b>
<b>8.7.</b>	<b>CONCLUSIONS .....</b>	<b>157</b>
8.7.1.	Theoretical contribution .....	157
8.7.2.	Pragmatic contribution.....	157
8.7.3.	Methodological contribution .....	158
<b>REFERENCES</b>	<b>.....</b>	<b>159</b>
<b>APPENDIX A: REVIEW OF RELEVANT OS METHODOLOGIES</b>	<b>.....</b>	<b>171</b>
<b>NORMATIVE APPROACHES</b>	<b>.....</b>	<b>172</b>
	Management Oversight Risk Tree (MORT) .....	172
	ACCIMAP.....	173
	Systemic Theoretic Process Analysis (STPA) .....	174
	System Failure Method (SFM) .....	176
<b>INTERPRETIVE APPROACHES</b>	<b>.....</b>	<b>178</b>
	Grounded Theory Methodology.....	178
	Historical Ethnography.....	180
	Causal Mapping.....	181
	Analysis of Vocabulary of Safety.....	183
	HRO Cases.....	185
	Disciplined Imagination.....	187
<b>SUMMARY</b>	<b>.....</b>	<b>189</b>
<b>APPENDIX B: PHILOSOPHICAL FOUNDATIONS</b>	<b>.....</b>	<b>192</b>
<b>POSITIVISM</b>	<b>.....</b>	<b>192</b>
<b>INTERPRETIVISM</b>	<b>.....</b>	<b>193</b>
<b>RATIONALE FOR THE CHOSEN PHILOSOPHICAL POSITION</b>	<b>.....</b>	<b>194</b>
<b>APPENDIX C: EXAMPLE OF AN NTSB SAFETY RECOMENDATION LETTER</b>	<b>.....</b>	<b>198</b>
<b>APPENDIX D: STUDY 1 DATA SET</b>	<b>.....</b>	<b>206</b>
<b>APPENDIX E: STUDY 2 INTERVIEW GUIDE</b>	<b>.....</b>	<b>209</b>
<b>APPENDIX F: STUDY 3 QUESTIONNAIRE</b>	<b>.....</b>	<b>212</b>
<b>APPENDIX G: LIST OF PUBLICATIONS PRODUCED FROM THIS RESEARCH</b>	<b>.....</b>	<b>218</b>

# List of Abbreviations

<b>ANSP</b>	Air Navigation Service Provider
<b>AOC</b>	Area of Change
<b>ARTS</b>	Automated Radar Terminal System
<b>AT</b>	Activity Theory
<b>ATM</b>	Air Traffic Management
<b>CA</b>	Conflict Alert
<b>CFIT</b>	Controlled Flight into Terrain
<b>CIM</b>	Contextual Implementation Model
<b>CPOE</b>	Computerised Provider Order Entry
<b>CSE</b>	Cognitive System Engineering
<b>DC</b>	Distributed Cognition
<b>EASA</b>	European Aviation Safety Agency
<b>FAA</b>	Federal Aviation Administration
<b>FMS</b>	Flight Management System
<b>GPWS</b>	Ground Proximity Warning System
<b>GTM</b>	Grounded Theory Methodology
<b>HAI</b>	Human Automation Interaction
<b>HCI</b>	Human Computer Interaction
<b>HFI</b>	Human Factors Integration
<b>HIT</b>	Healthcare IT
<b>HMI</b>	Human Machine Interface
<b>HRO</b>	High Reliability Organisation
<b>ICAO</b>	International Civil Aviation Organisation

<b>ISO</b>	International Organisation for Standardisation
<b>ISTA</b>	Interactive Socio-Technical Analysis
<b>IT</b>	Information Technology
<b>MAC</b>	Mid-Air Collision
<b>MORT</b>	Management Oversight and Risk Tree
<b>MSAW</b>	Minimum Safe Altitude Warning System
<b>NAT</b>	Normal Accident Theory
<b>NTSB</b>	National Transportation Safety Board
<b>OP</b>	Organisational Precursor
<b>OPHAII</b>	Organisational Precursors to Human Automation Interaction Issues
<b>OS</b>	Organisational Safety
<b>RA</b>	Resolution Advisory
<b>RC</b>	Requested Change
<b>SFM</b>	System Failure Method
<b>SME</b>	Subject Matter Expert
<b>SMS</b>	Safety Management System
<b>SOAM</b>	Safety Occurrence Analysis Methodology
<b>SPIN</b>	Safety Nets Performance Improvement Network
<b>STAMP</b>	Systemic Theoretic Accident Model and Process
<b>STCA</b>	Short Term Conflict Alert
<b>STPA</b>	Systemic-Theoretic Process Analysis
<b>UCD</b>	User Centred Design
<b>UK</b>	United Kingdom
<b>US</b>	United States
<b>VFR</b>	Visual Flight Rule

# List of Tables

Table 1. Overview of the three studies. (Specific details about data, data collection and data analysis procedures are reported in § 4.4 and 4.5.) .....	42
Table 2. Profile of Alphasky and Deltasky .....	54
Table 3. Study 1: list of participants interviewed.....	57
Table 4. Study 3: profile of the group of experts participating in the exercise. ....	59
Table 5. Summary of the coding methods used across the three studies.....	63
Table 6. Coding framework 1.1: example of a code. ....	65
Table 7. Framework 1.2: example of two codes.....	67
Table 8. Coding framework 2: example of the lower-level codes. ....	69
Table 9. Coding framework 2: example of case-categories (and related lower-level codes) for Alphasky and Deltasky.....	71
Table 10. Matrix of the changes requested by the NTSB to the FAA.....	87
Table 11. Safety recommendations containing the request to differentiate between the MSAW and CA aural alerts: for each recommendation, the table reports the NTSB safety letter in which the recommendation was made, and the ensuing letters exchanged by the NTSB and FAA. ....	88
Table 12. Coding framework 1.3: summary of the NTSB's and FAA's positions over NTSB Requested Change 1. ....	89
Table 13. Safety recommendation requesting FAA to modify existing regulations to make MSAW alarm transmission obligatory for controllers. The table reports the related NTSB safety letter and the ensuing letters exchanged between NTSB and FAA. ....	93
Table 14. Summary of the NTSB's and FAA's positions over RC7. ....	93
Table 15. Contrasting views regarding the role of the MSAW held by the NTSB and the FAA. References (included in brackets) refer back to the tables of Study 1 in which the particular category was identified.....	100

Table 16. The initial version of the OPHAI framework, as resulting from this study.....	102
Table 17. Summary of the categories identified in the Alphasky and Deltasky cases. (For examples of lower level codes and quotations see Table 9.) .....	105
Table 18. Organisational precursors to HAI issues resulting from the cross-case analysis. ....	115
Table 19. The updated version of OPHAI based on the findings of Study 2. While OP1 has remained unchanged, two more categories of precursors (OP2 and OP3) have been identified. ....	120
Table 20. Study 3: list of the corroboratory categories emerged in relation to the three categories of organisational precursors of the OPHAI framework. ....	122
Table 21. Evolution of OP1 as a result of Study 3: this category and its sub-items have remained unchanged. ....	125
Table 22. Initial (left box) and updated version (right box) of OP2 as resulting from Study 3.....	131
Table 23. Evolution of OP1 as a result of Study 3: this category has remained unchanged. ....	134
Table 24. The revised and expanded version of OPHAI, as resulting from Study 3. While OP1 and OP3 have been confirmed, OP2's subcategories have been refined and expanded.....	137
Table 25. Contribution made by the empirical studies to the OPHAI framework's development. ....	139
Table 26. Generalised version of the OPHAI framework. ....	140
Table 27. Summary table of relevant OS methodologies available in the OS. ....	171
Table 28. Control flaws leading to hazard (Leveson, 2002) .....	175
Table 29. Complete data retrieved from the NTSB and FAA databases during the initial phase of the study (see § 4.4.1.1). ....	206

# List of Figures

Figure 1. The pattern in organisational decision making leading to the Zeebrugge accident (Rasmussen & Svedung, 2000).....	25
Figure 2. The socio-technical control structure engaged in the control of risk (Source: Rasmussen & Svedung, 2000). .....	27
Figure 3. Theoretical gap addressed by this research.....	39
Figure 4. Study 3: organisational perspectives represented by the group of experts.....	60
Figure 5. Study 3: type of involvement of the participating experts in the safety net domain. ....	60
Figure 6. Type of safety nets the experts had direct experience with.....	61
Figure 7. Categories of concerns regarding the MSAW as identified by the NTSB over the period of analysis. ....	84

# Chapter 1.

## Introduction

Human automation interaction (HAI) can typically be found in safety-critical domains such as aviation, air traffic management, railways, and healthcare. Here, as part of their daily tasks, practitioners are provided with and are requested to use a variety of automated systems, such as information displays, decision support systems, and automated alarms. These systems augment human operators' ability to control the safety-critical processes for which they are responsible. This augmentation can occur in various ways, from registering data so that it is easily retrievable when needed; to analysing complex and dynamic information in order to support real time decision-making and action selection, or, in some cases, implementing autonomous system responses (Parasuraman, Sheridan, & Wickens, 2000). Ultimately, automation is an important asset for organisations operating in safety-critical domains, which can allow them to achieve increased productivity, quality of service, and safety.

However, although automation may succeed in realising these benefits, it can also have undesirable consequences at the human automation interface (Jones et al., 2011; Strong & Volkoff, 2010; Dekker, 2011; Woods, 2010; Degani, 2004). Automation may require operators to perform extra interaction tasks whose added value is not immediately clear to the practitioner. It may be accompanied with cluttered displays that make it difficult for practitioners to extract the salient information. Or, automation may flood the operational environment with nuisance alerts, i.e. alerts that, although generated by a system that works according to specifications, are perceived as irrelevant and distracting by the human operator. Overall, such undesirable consequences of automation, i.e., *issues in human automation interaction* (HAI), pose a relevant problem to the managers and administrators of complex, safety-critical organisations. Not only they may expose sharp end operators to confusion, uncertainty and frustration (Degani, 2004, p. 2); most importantly, they may compromise end-user acceptance, may increase costs, and may even lead to disastrous outcomes—such as harm, loss of life and of equipment.

For instance, the three recent commercial aviation accidents of Turkish Airlines Flight 1951 (Dutch Safety Board, 2010), Air France 447 (BEA, 2012), and Asana Airlines Flight 214 (NTSB, 2014) had as a common precursor the “automation paradox”, i.e., the tendency for pilots to become more and more reliant on complex automation to the extent that they lose their manual flying skills, a process that may result in the ineffective handling of emergencies when the automation fails and returns control to them. Other aviation accidents have been associated with other HAI issues, such as forgetting whether the autopilot is on or off, being confused about which localiser or beacon is in use, or about which way to measure altitude is in service (Perrow, 2011). By all means, the safety implications of HAI issues are not limited to aviation. In the marine domain, for instance, the introduction of the radar initially enabled vessels to move faster and more safely at night and during dense fog or storms; however, radar’s wide spread adoption has been linked to a rise in the so-called “radar-assisted” collisions—i.e., collisions caused by the tendency for multiple vessels to increase speed and perform sudden changes in direction whilst ignoring the possibility that other vessels can do the same, because provided with the same technology (Perrow, 2011). In the healthcare domain, issues with information presentation in systems such as computerised physician order entry may actually heighten—rather than reduce—the risk of medication error and resultant harm to patients (Koppel et al., 2005; Nebeker, Hoffman, Weir, Bennett, & Hurdle, 2005). Overall, these instances show that while automation has the ability to bring safety and efficiency improvements in virtually any domain, its potential to negatively impact system



safety should not be neglected. HAI issues, if not adequately handled, can make complex safety-critical systems vulnerable in new and unintended ways (Woods, 2010; Perry, Wears, & Cook, 2005).

### **1.1. RESEARCH GAP**

---

HAI issues are not a recent problem. There are at least two perspectives, or areas of study, that are relevant to these issues. The first is human computer interaction (HCI). As it will be discussed in chapter 2, HCI scholars have extensively problematized the interaction that may unfold between the human and the machine during the use of the latter. Theoretical perspectives such as human information processing, activity theory, distributed cognition, computer supported collaborative work, and cognitive system engineering have produced profound insights into the cognitive, social and contextual dimensions that may influence HAI. In particular, they have provided theoretical models useful to characterise HAI, either in laboratory settings and/or in complex dynamic environments; to identify what may go wrong in the use of technology, namely how and what HAI issues may occur; and to understand how such issues can be removed by means of corrective design changes. The second perspective, here named as the system lifecycle perspective, includes process models aimed at improving the fit between automation and the intended user. User centred design is perhaps the most notorious of these models, and it has been developed in order to ensure that the needs of the end user are adequately considered during the system development process. Other relevant approaches are human factors integration and system safety. When applied to the development of automation, these approaches provide a useful means to properly capture, identify HAI issues, and translate them into implemented user, human factors, or safety requirements.

However, while both the HCI and the system lifecycle perspectives have provided useful insights into HAI issues, they neglect the organisational sources of the problem. Yet this is an important question to address, for at least two reasons: a pragmatic, and a theoretical one. First, the increasing pervasiveness of HAI in safety critical domains is paralleled by the occurrence of accidents in which HAI issues play a contributory role, as mentioned earlier. This urges researchers to deepen our understanding of the sources of these issues, and develop additional explanations to those provided by the current HCI and the lifecycle perspectives.

Second, organisational precursors have been extensively investigated in the organisational safety (OS) area. There, it has been noted that complex, safety-critical domains can be seen as a wedge, which consists of a *sharp end*, i.e. the level at which automation is used by the human operator; and a *blunt end*, the level at which relevant decision makers such as engineers, management, regulators, manufacturers and the like can be found (e.g. Dekker, 2012; Flin, 2008; Reason, 1997).

OS scholars have long highlighted that there exist specific organisational precursors found at the latter level that can predispose an organisation to successful safety performances, or on the contrary, can lead to major accidents and disasters (e.g., Leveson, 2012; Perrow, 2011; Woods, 2005; Hollnagel, 2004; Reason, 1997; Snook, 2002; Rasmussen & Svedung, 2000; Turner & Pidgeon, 1997; Vaughan, 1997). Such precursors include for instance blind spots in organisational decision making; presence of strong managerial pressures for cost effectiveness and productivity, rather than safety; failure to enforce safety constraints at various organisational levels, from government and regulators, down to company management and operations; tendency to gradually tolerate more and more level of risks, thus resulting on the organisation accepting more risks that it can safely handle. These precursors are known to predispose the organisation to an incremental descent into failure, a pattern also known as organisational drift into failure (e.g., Dekker, 2012; Farjoun, 2005; Woods, 2005).

Other scholars have highlighted the precursors that can resist this drift (Roe & Schulman, 2008; Schulman, Roe, Eeten, & Bruijne, 2004; Porte & Consolini, 1998; LaPorte, 1988). These include, for instance, the availability of a top management committed to safety; of an organisational environment that promotes organisational safety learning; of organisation members committed to continuously challenge existing beliefs about the safety of the organisation.

Overall, the identification of these organisational precursors has proved immensely helpful for understanding the organisational sources of failure, and for devising safety improvement measures targeting organisational and managerial levels, rather than the level of operations (the sharp-end). (This is in fact the level targeted by classic safety approaches, which notably focus on human error and technical malfunctions.) However, OS scholars have investigated organisational precursors in relation to accidents and disasters, and not problematic human automation interaction. Therefore, it remains to be explored whether similar types of precursors exist in relation to HAI issues.

## 1.2. AIM AND OBJECTIVES

---

The aim of the present research is to advance current knowledge of how organisations operating in safety-critical domains manage (or mismanage) the HAI issues accompanying the automated systems they provide to their sharp end operators. The specific objective is to explore and identify the organisational precursors to HAI issues found at the blunt end of safety-critical domains. In particular, the objective is to develop a qualitative framework of the organisational precursors to HAI issues (OPHAI). Consistently with the contributions from the OS area, organisational precursors include those organizational dynamics and conditions that, although distant in time and space from operations, may actually influence the way in which HAI issues are successfully addressed and mitigated, or remain unaddressed in the organisation—thus remaining visible to the user of the technology.

It is anticipated that this research has a strong qualitative (or interpretive) orientation. In particular, the OPHAI framework is developed in a bottom-up manner from the data collected across three successive qualitative studies: two qualitative, historical case studies followed by a third subject matter expert study. Study 1 delivers the first theoretical category of organisational precursor of the framework. While identifying added support for this category, Study 2 identifies two more theoretical categories. Study 3 expands and refines these three categories in order to deliver the final version of the OPHAI framework. The focus of the research is mainly on the organisational context of implementation and improvement of an automated system from the air traffic management domain, the Minimum Safe Altitude Warning system, or MSAW. For this reason, the research design can be considered as a derivative of the single case study design, frequently used in the OS area. The data considered along the OPHAI framework's development includes qualitative data such as documentary data; semi-structured interviews; field notes and observations; and a qualitative questionnaire.

## 1.3. THESIS STRUCTURE

---

The following two chapters review literature that establishes the context of and the theoretical motivation for this research. Chapter 2 reviews the two main current theoretical perspectives on HAI issues, namely the HCI and the system lifecycle perspectives. Chapter 3 introduces the third, alternative theoretical perspective, namely the OS perspective. The end of the chapter relates this perspective to the HCI and the system lifecycle perspectives reviewed in chapter 2 in order to consolidate the research gap the present study addresses.

Chapter 4 describes the research strategy. More specifically, this chapter describes the overall research plan by which the research objective is achieved, and the lower-level methodological aspects of data collection and analysis used within each study.

Chapters 5 to 7 present the results of the three qualitative studies. Chapter 6 describes the results of Study 1; chapter 7 describes those of Study 2; and chapter 8 describes those of Study 3. Before reading these chapters, the reader should note that only minimal reference is made to existing literature in each study. It is the task of the final chapter, chapter 8, to discuss the results of this thesis in relation to existing relevant literature. The chapter also concludes the thesis by listing the theoretical, pragmatic, and methodological contributions it makes.

# Chapter 2.

## **HAI ISSUES: CURRENT THEORETICAL PERSPECTIVES**

### **2.1. CHAPTER INTRODUCTION**

---

This chapter reviews the main current theoretical perspectives used to understand and deal with issues in human automation interaction (HAI). These perspectives have been grouped into two major clusters:

- Perspectives looking at the analysis of the human machine system unit. This cluster includes classic models from the Human Computer Interaction area that have been used to understand human technology interaction in safety-critical domains, such as human information processing, distributed cognition, activity theory, computer supported collaborative work, and cognitive system engineering.
- System lifecycle perspectives. This cluster includes process models and frameworks that promote the integration of user, human factors and safety requirements along the lifecycle of automated systems. Three main perspectives have been included in the review: user centred design, human factors integration, and system safety.

The objective of the chapter is to highlight the contribution that each perspective has brought to the problem of HAI issues, and, in particular, to understand how such perspectives consider the organisational precursors to such issues.

## 2.2. THE HUMAN COMPUTER INTERACTION PERSPECTIVE

---

Human Computer Interaction (HCI) is the classic perspective to bring over the problem of HAI issues. It is common among Human Factors and Human Computer Interaction scientists and practitioners with an interest in automation and the problems of poor fit which can arise from the use of automated systems. This perspective includes those theoretical approaches that have defined their unit of analysis as composed of (i) the human (be it a single operator or a team), (ii) its (or their) automated system(s), and (iii) the interaction that develops between (i) and (ii). As it will be described later on this chapter, differences can be observed across these perspectives depending on their consideration of the context surrounding the human machine unit; however, the fundamental idea remains that a proper conceptualisation of this human-machine system unit is what supports the identification of potential interaction design flaws and mismatches, which in turn can be designed out—hence achieving an increased human-centeredness. Conversely, a limited conceptualisation of this unit would lead to a poor fit between the human and the system.

The remainder of this section will review the most notable works belonging to the HCI perspective, and it will conclude with a critique of the perspective regarding its ability to consider the organisational precursors that this research is set to examine.

### 2.2.1. Human information processing

From a human information processing perspective, HAI issues reflect a basic mismatch between the features of a tool and the cognitive processing abilities of individuals. The human information processing paradigm includes theoretical models that are grounded on the metaphor of the human operator as an information processing system. From this angle, human cognition is seen as composed of different processing structures which process incoming information before a response is formulated (Harris, 2012, p. 20).

Human information processing (HIP) has the merit of having stimulated useful laboratory research addressing relevant constructs useful for characterising HAI issues. For instance, HAI issues may be characterised in terms of an excessive *trust* in automation. Notably, as their trust in automation grows, human operators tend to delegate an increasing amount of tasks to their automated systems (Parasuraman & Riley, 1997). The side effects of such increased trust in automation is that they spend less time on manual control during normal operations; consequently, their control and problem solving skills deteriorate over time because they are not used as continuously as they would in situations with manual control. As a result, the skill degradation associated with the prolonged reliance on automation during normal operations may result in a decreased ability to respond to

abnormal situations. In these situations control might shift back to human operator who might lack the expertise to execute a timely and appropriate response. Actually, lacking critical knowledge, human operators might even aggravate the situation by committing a fatal error (Bainbridge, 1983).

In an opposing case, trust can be compromised by automated systems that are found to work unreliably, so that operators may start to ignore them or switch them off. This is the classic problem when practitioners are exposed to an excessive rate of false or *nuisance alerts*, i.e., alarms that do not correspond to the presence of a real danger. Usually generated by automated alarms, nuisance alerts are annoying and distracting to human operators as they produce frequent interruptions in the flow of activity. As a consequence, operators might become increasingly desensitised as they lose trust in their alarms (Keller, Diefes, Graham, Meyers, & Pelczarski, 2011), something that may lead to slow reactions to true alarms, or even to disregard alarms entirely (Grounds & Ensing, 2000).

While trust is just one useful construct that arises from the use of the HIP paradigm, others exist too—e.g., cognitive workload, situation awareness, fatigue, and reliance in automation (Harris, 2012). These constructs are valuable as they provide scholars with a language to describe HAI issues in terms of cause effects relationships, to develop predictions about user behaviour, and to carry out laboratory evaluations of different interactive systems, to investigate their adequacy for supporting various cognitive tasks (Rogers, 2012, p. 25). However, one of the main limitation of the human information processing paradigm lies essentially on its experimental nature: it assumes that human cognition can be abstracted and investigated in laboratory settings, in isolation from the naturalistic context where HAIs occur (Hollnagel & Woods, 2005). As a reaction to this position, other alternative approaches have been developed by scholars who took the use of technology in context as their focal concern. These alternative approaches are reviewed in the next sections.

### **2.2.2. Distributed cognition**

Still a theory of cognition, distributed cognition (Hollan, Hutchins, & Kirsh, 2000) differs from the human information processing paradigm because it sees cognitive activities as not limited to the individual, but as a property of humans and the context in which they are located. Distributed cognition (DC) sees cognitive processes as distributed across individuals, artefacts (both paper-based and computer-based) and the environment. These are, according to DC, the components of distributed cognitive systems.

A classic example of a distributed cognitive system is the aircraft cockpit. In analysing this kind of system, distributed cognition is mainly concerned with information and its propagation through the components of the distributed cognitive system. More specifically, DC analyses the creation, transformation, and propagation of representational states in the system in order to find possible breakdowns. Breakdowns may occur because information does not propagate effectively from one representational state to another (Rogers, 2012, p. 39). For instance, this may be due to a mismatch between the information that is provided by an information artefact and the expectations of the human about that information.

Distributed Cognition has not highlighted new types of HAI issues specifically. It has mainly provided practitioners with a framework useful for describing in depth the normal functioning of a work situation, and identify possible sources of breakdowns within it (e.g., Rajkomar & Blandford, 2012).

### **2.2.3. Activity theory**

An alternative perspective to DC is Activity theory (AT). Similar to the former, AT also looks beyond individual cognitive processing. However, while DC focuses on distributed cognitive systems, AT has drawn its unit of analysis around the concept of object-oriented, collective, and culturally mediated human activity. Notably, such an unit of analysis has found expression in the popular Engestrom's model of AT, a model whose components include the set of goals, motivations, communities, rules, division of labour, artefacts and computer systems that accompany situated human practice (Engestrom, 2000). AT promotes the analysis of the multiple and unfolding relationships between these components in order to find *contradictions*, or inconsistencies between them. For instance, contradictions may arise when the highly hierarchical division of labour of a given work setting is not supported by the model of work that has been built into a new device that has to be used in that setting (Bonneau, 2013).

AT, similar to DC, has the merit of having provided a qualitative framework for developing contextual analysis of human activity in complex socio-technical systems, making inferences across interactions, findings patterns, and describing them. Similarly to DC, this approach forces the analyst to consider the broader context in which human activity is located and the technology used, although it arguably promotes a deeper understanding of such a context, considering its strong emphasis on the motivational and historical components of human activity.



#### **2.2.4. Computer supported collaborative work**

Computer supported collaborative work (CSCW) is an area of study that has focused on investigating how technology supports collaborative work in context. Thus CSCW focuses on the study of local collaborative practices found across the members of both co-located and (massively) distributed teams.

Usually grounded on ethnographic studies of work, CSCW has highlighted different dynamics that are involved in the collaborative activities of operators of complex settings. One important dynamic is that of shared interpretation (Bannon, 2000; Bannon & Bødker, 1997; Reddy, Dourish, & Pratt, 2001)—i.e., how different operators can maintain a sufficiently coherent interpretation of the activity of other co-workers, either co-located, or dislocated, functional to do the work. Such a shared interpretation provides the context for own activity, in fact. An essential means in sustaining such an interpretation is provided, for instance, by coordination artefacts—digital and physical artefacts that are used to keep operators up-to-date about key information, and to simplify coordination among individuals (e.g., Redaelli & Carassa, 2015; Bardram & Bossen, 2005). For instance, the rack of flight progress strips—the classic example of a coordination artefact from the air traffic control domain (Hughes et al., 1992; Harper and Hughes, 1993; Fields et al., 1998; Berndtsson and Normark, 1999)—provides controllers with a very compact and at-a-glance representation on the status of traffic flow in a sector, something that supports individuals when carrying out their tasks, but it also supports essential cooperative work among different controllers (Fields, Amaldi, & Tassi, 2005).

Also, other components of cooperation involve glancing at what others are doing, overhearing on-going conversations of other co-workers, and establishing conditions for tacit coordination. Subtle, ingenious, and practical strategies such as these have been highlighted since the 90s by the ethnographies of work of a group of British sociologists from Lancaster University, who studied in depth the collaborative work as found in air traffic control (Harper, Hughes, & Shapiro, 1991, 1989) and ambulance control centres, and the London underground (Heath & Luff, 1991).

Overall, CSCW works are important because they warn that technology, if not carefully fitted to the context of use, might inadvertently affect not just individual practices, but also collaborative ones. In fact, the failure to comprehend a situated context of work may result in the technology constraining the existing collaborative practice in unintended ways.

### 2.2.5. Cognitive system engineering

Cognitive system engineering (CSE) is a multidisciplinary approach to the analysis, design, and evaluation of complex human machine systems (Hollnagel & Woods, 2005; Rasmussen, Pejtersen, & Goodstein, 1994). Influenced by many disciplines, such as system engineering, human factors, cognitive and ecological psychology, this approach is mainly concerned with supporting cognitive performances in complex safety-critical domains such as healthcare, aviation, air traffic control, and nuclear power plant. One of the main differences compared to HIP is the system theory foundation of CSE: rather than viewing human-machine interaction as a decomposable in a mechanistic sense, this approach maintains a primary concern with the synergistic functioning of joint cognitive systems—i.e., interactive teams of human and technology (Hollnagel & Woods, 2005)—that are in control of safety critical systems.

Works from CSE have been effective in identifying severe HAI issues that have led to safety-critical outcomes across different domains. For instance, *automation surprises* (Sarter, Woods, & Billings, 1997; Woods & Sarter, 2000) capture those situations where human operators are surprised by a system that is operating in a mode that is different from the one expected by the user. This problem arises (i) because automation has multiple operating modes and (ii) because it makes it possible to program long and complex sequences of actions whose execution might not be entirely intelligible to the humans who are in control (Norman, 1990). As a result, human operators might lose track of automation operating modes and develop a set of expectations that is inconsistent with what the automation has actually been set up to do (Woods, 2010, p. 148).

The severity of automation surprises can be appreciated, for instance, in relation to the flight management system (FMS) available on commercial airliners. Incidents and accidents relating to FMS use have shown that pilots might be surprised when they realise that their aircraft's behaviour does not match their expectations (Woods, 2010, p. 148). Another classic HAI issue is the problem of *clumsy automation*. First noted by Wiener (1989), this issue captures the fact that automation, while reducing workload during the longest and less demanding phases of flight, i.e., en route, actually increases operator load during the most demanding phases of the flight, i.e., take-off and landing. Here, the pilots might be exposed to additional attentional, communicative, and coordinative demands—demands that in turn may create opportunities for novel types of errors and novel paths to system failure (Woods, 1996, p. 2).

Also, automation may come with additional knowledge demands, as operators may be required to learn and familiarise themselves with a large set of (novel) functions and know

when to activate them. This requires operators to dedicate more time to remember “input models, understanding display readings, setting up and initializing devices, configuration controls and operating sequences” (Woods, 2010, p. 145). So, the complexity of automation might lead operators to alter or dismiss some of the system’s functionalities to achieve a level of simplicity and ease of use that is more adequate to their operations.

Woods et al. (2010) have reported an example where clinicians set up their devices in order to minimise their need to interact with the new technology during high-tempo periods. This occurred despite the fact that the practitioners’ configurations neutralised many of the new systems’ expected advantages—e.g., the flexibility to perform a wider range of different kinds of data manipulation (Woods et al. 2010, p. 192).

Adjustments such as these are not surprising considering the adaptive nature of human behaviour in safety-critical systems. Humans are not passive receivers of automated tools. HAI is situated within an operational context where humans constantly trade-off between alternative courses of action in order to meet the multiple and conflicting demands of their job in a cost-effective way, considering available (scarce) resources, existing institutional objectives, severe productivity and temporal pressures (Hollnagel, 2012b; Woods, Dekker, Cook, Johannesen, & Sarter, 2010). As a consequence, it is not surprising that the constant search for improved cost-effectiveness pushes operators systematically to turn to informal usages (Wright & McCarthy, 2003), i.e., usages that differ from that prescribed by system developers, management and regulators. Note that informal usages do not necessarily lead to adverse consequences, provided that an organisation establishes appropriate controls to detect them and coordinate them across their different units. Otherwise, the risk is that of having local adaptations that while looking perfectly reasonable from the perspective of the local user, may actually compromise the integrity of overall operations (Rasmussen & Svedung, 2000).

### **2.2.6. Critique of the HCI perspective**

The previous sections have reviewed various theoretical models that belong to the HCI perspective in order to see the value that these models have brought to the understanding of HAI issues. From the review, it is evident that these approaches provide both researchers and practitioners with a set of concepts and a language that are useful for characterising different types of HAI issues. In particular, the main value of these perspectives lies in their diagnostic power: they can help researchers and practitioners to trigger diagnostic questions about a given human machine system—e.g., is the system reliable enough to enable users to trust it? Does the system provide sufficient feedback to its user to inform her/him about its behaviour? Is there a risk of mode confusion?—and about the environment that shapes the usage of the system—e.g., is the temporal demand compatible with the time frame that is required to operate a system? Are existing collaborative patterns adequately supported by automation? Questions such as these are immensely helpful when analysing HAI problems, as they lay the foundations for further formal and informal evaluations, as well as corrective design changes.

However, the HCI perspective alone has a limited explanatory power with regards to the appreciation of the organisational precursors to HAI issues. In fact, this perspective does not conceptualise the organisational sources of the problem. For instance questions such as how biased regulator, managerial or engineering decisions may lead to HAI issues, or how organisations may actually (mis) handle HAI issues cannot be addressed by the HCI perspective alone, as these kinds of questions do not pertain to the unit of analysis of this perspective.

### **2.3. SYSTEM LIFECYCLE PERSPECTIVE**

---

While the perspectives reviewed in the previous section focus on the interaction between humans and automation, this section reviews works that draw the boundaries of their unit of analysis around the broader engineering lifecycle of an automated system. Analogously to the human lifecycle phases of birth, childhood, death, and burial the lifecycle process of an automated system includes five major phases at least (Ericson, 2005): concept definition; development; implementation; operation; and decommissioning. These phases correspond to subsequent system transformations that are linked by predefined input output relationships. They are considered to apply to both hardware and software systems, except that for the latter the development and test phase exists in many variations, such as the Waterfall model, the Spiral Model, the IEEE V model and the like.

Within such an engineering lifecycle or process, HAI issues can be seen as process failures. They are seen, in fact, as a reflection of the inability of development teams to properly capture human factors and safety requirements and realise them into the implemented system. This view of failure originates from a comparison of development practices that happened in real life with prescriptive process views as mandated by the user centred design, human factors integration or safety perspectives. Notably, these approaches prescribe the use of different principles and methods along the system development lifecycle to ensure that user, human factors, and safety requirements are actually identified and built into the final system. These three approaches are reviewed hereafter.

#### **2.3.1. User centred design**

User centred design (UCD) is a design philosophy that emphasises the continuous involvement of end-users throughout the different phases of the system development lifecycle. Such continuous involvement aims to ensure that the design of novel systems is actually driven by end-user expectations, desires and needs as captured in user requirements—ultimately enhancing human-system interaction. The User Centred Design process may come in different variations, and usually it involves, along the system development lifecycle, the use of different methodologies and techniques such as ethnographic studies, contextual inquiry, prototype testing, and affinity diagrams. One useful UCD reference has been provided by the International Organisation for Standardisation (ISO) (2010), which has codified the six principles that allow the verification of whether a design process is user centred: (1) the design should be based on the explicit understanding of the users, tasks, and environments; (2) the users should be involved throughout design and development; (3) the design should be driven and refined by user centred evaluations; (4) the process should be iterative; (5) the design

should address the whole user experience; (6) the design teams should include multidisciplinary teams and perspectives.

### **2.3.2. Human factors integration**

A framework adopted by defence and hazardous industries, human factors integration (HFI) is aimed at coordinating and assuring the integration of human factors methods within the engineering lifecycle of a system. Classic areas targeted by HFI include interface design and workplace layout, trust and acceptance of the system, longer term planning and staffing, skill changes and training, human error and recovery, and so on (e.g., Shorrock, Woldring, & Hughes, 2004; Widdowson & Carr, 2002). Specific HFI methods useful during the development and evaluation of automated system include hierarchical task analysis, cognitive task analysis, human reliability assessment, cognitive walkthroughs, and human-in-the-loop simulations (Lowe, 2008). Note that none of these methods has been defined for addressing HAI issues specifically, as human factors integration is concerned with all of the human issues that may originate from the introduction of a novel system—such as manpower, personnel, training, and health and hazard—hence it is not limited to HAI issues.

### **2.3.3. System safety**

System safety can be defined as the process of managing the safety risks that arise along the lifecycle of a new system. At its core the approach holds the idea that safety is a system property that can be built into the system since the early conception phase. Indeed, the experience with complex system failures has shown that this approach is much more cost effective than adding safety improvements in the aftermath of accidents and disasters. The ideal objective of this approach would be to eliminate any potential hazard that could be introduced by a new system; however, this is not always feasible, especially when dealing with complex systems such as nuclear power plants, weapon systems, and aircraft (Ericson, 2005). Therefore, system safety strives to reduce risk to an acceptable level. It does so by focusing on the systematic, forward looking identification and elimination or mitigation of safety hazards throughout the system lifecycle (Roland & Moriarty, 1990). For instance, this is achieved by carrying out specific safety assurance tasks along system lifecycle phases, such as a functional hazard assessment during the initial system definition phase, preliminary safety assessment during the system design phase, and system safety assessment during implementation and transfer to operation (Kirwan, 2007; EUROCONTROL, 2001).

Usually this process can be complemented by the use of a safety case, and it is executed in the context of a safety management system. The former identifies a formal document

that contains a structured safety argument that documents the evidences, intermediate conclusions, and assumptions to demonstrate that a given design solution meets the agreed safety levels. The latter usually identifies a coherent managerial framework that is aimed at establishing different safety functions within an organisation—such as risk assessment, incident reporting and investigation, safety monitoring—and an overall culture functional to the achievement of good safety performances (ICAO, 2013; Stolzer, Halford, & Goglia, 2008). Essentially, the safety case and the safety management system respond to the regulatory requirement to demonstrate that a system is safe to operate.

Also, system safety, similar to human factors integration, has the capability to address various kinds of safety issues, rather than just HAI issues alone. It usually applies to different kinds of systems and contexts, not necessarily human controlled, and, compared to human factors integration, it also addresses technical and engineering issues rather than solely human-related ones. In fact, the focus of system safety is not, on improving the overall fit between the human and the system, but to achieve an acceptable level of overall system safety. Regardless of these considerations, the two approaches are not mutually exclusive but have the potential for complementation (Lowe, 2008).

### **2.3.4. Critique of the system lifecycle perspective**

Overall, UCD, HFI, and SMS provide useful means for improving the fit between the human and the technology. They can provide a process explanation of HAI issues: these models provide some useful potential references against which to carry out retrospective reviews of past and current development practices, identify process limitations in current human factors, and safety assurance practices.

For instance, the misfit between human and their equipment has notably been associated to the late involvement of human factors specialists along the system development process—a late involvement that happens only when the most relevant design decisions have been taken already, and when there are limited margins for analysis and improvement (Cardosi, 1998; Kirwan, et al., 1997; Leveson et al., 2001). Other process limitations include, for instance, the failure to capture relevant safety issues; incomplete safety case evidence resulting from the limited time allocated to HMI evaluation; tendencies to rely on contractors who will sign off safety tasks without knowing how these will relate to other aspects of system safety; pressures when classifying safety issues, such as de-rating a safety issue as a workload issues (Humphreys, Kirwan, & Ternov, 2006). Others have added to this list the tendency for development teams to fail to include automated system in the scope of safety assurance activities, based on the biased assumption that such systems play a supportive, advisory role only—but cannot cause or contribute to an accident (Sandom, 2009; Sujan, 2001).

These types of insights are useful because they allow researchers and practitioners to remove bottlenecks and limitations in current or/and past user centred design, human factors, and safety assurance practices from forthcoming development practices. At the same time, it can be noted that the same types of insights fail to reach the deeper managerial and organisational sources of the problem. For instance, linking HAI issues to the late involvement of human factors personnel does not shed light on the organisational sources of the problem: how and why do organisations that operate safety-critical systems may come to tolerate poor human factors integration or safety assurance practices when introducing safety-critical automation? Or, how do organisations direct their resources and efforts towards the adoption of the three approaches discussed in this section in order to ensure that there is an adequate fit between automation and its users?

Overall, questions such as these point at the need to consider the wider organisational context in which automation is deployed and used. Being able to shed light on this context would provide an understanding of the possible organisational sources of HAI issues, organisational sources that in turn could become targets for improvements that are complementary to those that have been identified within the HCI and the system lifecycle perspectives.



## 2.4. CHAPTER CONCLUSIONS

---

This chapter has reviewed the two main perspectives that are available today to deal with the problem of HAI issues: the HCI and the system lifecycle perspectives. The chapter has concluded that while both perspectives have brought important insights into the problem of HAI issues, they also neglect the organisational precursors to these issues. In fact, these precursors fall outside of the scope of such perspectives.

The classic theoretical models available under the HCI perspective have the merit of providing a means to conceptualise and identify HAI issues, which in turn can be removed from the system by corrective design changes. By electing the human machine unit (and its situated context in some cases) as its unit of analysis, this perspective ignores the organisational factors that influence how HAI issues, which occur within such a unit, are managed.

Compared to the HCI perspective, the system lifecycle perspective provides a somehow enlarged view. This perspective includes process models that mandate, along the engineering lifecycle of a system, the forward looking consideration of the end-user, human factors and system safety aspects—a consideration that can increase the human centeredness, efficiency and safety of human machine systems. Therefore, such perspective can trace the occurrence of HAI issues to process failure—the failure to implement basic user centred design, human factors integration and system safety processes and principles along the system development lifecycle of a system. By doing so, this perspective can shed light on process aspects that are not touched by the human machine system perspective; however, it still fails to reach the higher-level organisational precursors of HAI issues. The engineering lifecycle of a system does not happen in a vacuum, in fact, but it is located into a specific organisational context.

These considerations call for a third, enlarged, theoretical perspective—a perspective able to appreciate the organisational precursors to HAI issues. Arguably, appreciating such precursors would provide opportunities for improvements that current perspectives are unable to identify. Such a third perspective is introduced in the next chapter.

# Chapter 3.

## **THE ORGANISATIONAL SAFETY PERSPECTIVE**

### **3.1. CHAPTER INTRODUCTION**

---

The previous chapter suggested that current approaches to human automation interaction (HAI) issues do not necessarily consider the deeper organisational and managerial sources of the problem. This chapter outlines a third perspective, namely the organisational safety (OS) perspective, which incorporates the organisational precursors to accidents and system safety.

Section 3.2 provides an introduction into the OS perspective and the major pattern of failure it contemplates, i.e., organisational drift into failure. Subsequently, sections 3.3 and 3.4 review fundamental theoretical models to respectively explain how and why organisations may drift into failure or how they may avoid this. Finally, section 3.5 links the OS perspective with the perspectives presented in the previous chapter, in order to define the research gap of this research.

### 3.2. THE ORGANISATIONAL SAFETY PERSPECTIVE AND ORGANISATIONAL DRIFT INTO FAILURE

---

At the core of the OS perspective lies the idea that complex safety-critical systems can be conceived as a wedge formed by a sharp end and a blunt end (e.g. Dekker, 2012; Flin, 2008; Reason, 1997):

1. The *sharp end* includes those operators, such as pilots, air traffic controllers, and healthcare professionals, who—while exploiting a variety of resources, including automated systems—are directly responsible for controlling the highly hazardous process they are responsible for, e.g., flying an aircraft, instructing aircraft in order to establish safe separations among them, or provide medical assistance to patients, respectively. Due to their proximity to such processes, the activity of these operators is usually characterised by the need to make real time decisions in a highly dynamic environment;
2. The *blunt end* includes those organisational actors, such as engineers, managers, policy makers and regulators, who constrain, put pressures upon and ultimately influence the activity of the operators at the sharp end. Humans at the blunt end are not in contact with the hazardous process to be controlled. They lack, in fact, temporal and spatial proximity with the operational context, and because of this they have also been referred to as “behind-the-scenes” actors—behind the scenes in that their activity is not directly visible from the perspective of the operators at the sharp end (Balca, Doyle-Waters, Lecznarowicz, & FitzGerald, 2007).

Notably, while the sharp end has been the area of analysis and intervention in classical safety approaches, which notably have focused on addressing technical malfunctions and human errors, evidence from accidents occurring in complex safety-critical systems has suggested that the blunt end should also be considered a valuable area of analysis and safety improvement (e.g., Dekker, 2011; Catino, 2006). As it will be discussed throughout this chapter, accidents (here defined as any catastrophic occurrence, potentially leading to fatal or serious injuries, damage or loss of equipment) do not originate solely from individual errors or technical failures at the sharp end, but their roots can be traced to organisational conditions at the blunt end.

This idea is not new. The International Civil Aviation Organisation (ICAO) (ICAO, 2013) explains that safety has evolved along three eras: the technical era, the human factors era, and the organisational era. Since the early 1900s, safety and its counter side—accidents—were considered, essentially, as induced by technological failures. Therefore,

great emphasis was placed on diagnosing problems with technical systems. Starting in the early 1960s, the scope of safety science was extended to include consideration of the human contributors to accidents. This extended scope led to the consideration of aspects of human performance that were not considered during the earlier technical era. One limit of the human factors era was the predominant focus on the individual (and his/her errors), with limited consideration of the surrounding organisational context. It was only in the early 1990s, due to the seminal work of social and organisational scientists with an interest in safety, that attention was drawn to consideration of the surrounding organisational context in which humans operate and in which accidents develop. Since the beginning of the organisational era, accidents have also been viewed as the result of organisational dynamics and conditions at the blunt end, rather than purely human or technical failures at the sharp end. Accidents, in other words, have been considered to be an organisational phenomenon; for this reason, the notion of “organisational accident” was introduced.

The notion of organisational accident reminds us that failure is rarely the result of isolated human or technical failures at the sharp end. In addition to operations, there is a surrounding organisational context that may provide the foundation for a catastrophic failure. The analysis of accidents such as Three Mile Island, Bhopal, Chernobyl, Challenger and Columbia has shown us that accidents can originate from organisational conditions such as biases and blind spots in organisational decision making, impaired organisational safety learning, and from strong political, managerial, and economic pressures for productivity in spite of safety objectives (Dekker, 2011, p. 122). Organisational conditions such as these, apparently unrelated to operations, may set the organisation on a slow descent into failure. This idea, of organisational drift into failure, receives support from a variety of models from the OS area. The remainder of this section will review the most important of these models, in order to highlight the specific traits of organisational drift.

### **3.2.1. A by-product of normal organisational and administrative activity**

Patterns of organisational drift result from the “imperfect” or biased decisions made by various organisation members over time, while engaged in their daily work, rather than from something exceptional or extraordinary act. Biases in everyday decision-making in complex, safety-critical organisations have their foundations in the *local rationality principle* and the presence of tensions for *cost-effectiveness*. Elaborated by Herbert Simon (1957), the former suggests that organisation members or actors do not behave as perfectly rational decision makers, i.e., decision makers who choose the best course of action after having exhaustively enumerated all possible options. Rather than behaving in this way, members ground their behaviour on their local understanding, i.e., their

### Chapter 3. The organisational safety perspective

understanding of the work situation as seen from their unique “position” in the organisation hierarchy, given existing information and attentional resources, demands, goals and context. The presence of limited local rationalities can be seen as one of the side effect of the division of labour, which in large organisations usually produces complex organisational structures, characterised by several hierarchical levels and organisational boundaries. While functional to the achievement of the institutionally relevant objectives, these structures and boundaries may also obstruct the flow of information across different units, departments, and teams in ways that are dysfunctional to effective safety management.

When translated to HAI issues, these considerations prompt exploration into whether it is possible to trace HAI issues to past imperfect managerial and engineering decisions—decisions that, while appearing rationale at the time they were made, considering the specific situated perspective of a given actor, may have actually ignored the HAI issues accompanying a given automated system, or may have constrained effective responses to such issues.

Furthermore, decisions are usually imperfect as they are made under strong pressures for cost-effectiveness. Complex, safety-critical organisations need to continuously balance efficiency versus safety goals (Hollnagel, 2012b; Vaughan, 2009; Cook, Nemeth, & Dekker, 2008; Karen Marais & Saleh, 2008; Woods, 2005; Rasmussen & Svedung, 2000). Organisational efficiency implies directing organisational resources towards the achievement of (short term) economic goals. This may occur, for instance, by increasing production rates, capacity utilisation, minimising slack in the system, reducing personnel and trimming costs (Shrivastava, 1994). In contrast, organisational thoroughness implies directing organisational resources towards the achievement of (long term) quality and safety goals. This may occur, for instance, by conducting hazard analysis methods during system design and development, by performing root cause analysis to learn from past incidents and accidents, and by implementing safety case regimes, safety management system and safety culture programs to improve safety oversight. When regarding the management of HAI issues specifically, this implies for instance directing adequate resources towards the adoption of the already mentioned user centred design and human factors integration approaches (§ 2.3).

The need to balance between efficiency and safety permeates through all organisational levels, from sharp end practitioners to top management. For instance, when facing a strong demand for productivity goals, sharp end operators will reduce thoroughness until productivity goals are met; conversely, when facing strong safety goals, operators will

### Chapter 3. The organisational safety perspective

reduce efficiency until safety goals are met (Hollnagel, 2012b). Similarly, top management, in the absence of accidents, may direct most of their resources—people, funds, expertise, attention and equipment—to attain productivity goals, thus meeting existing short term demands for cost effectiveness while sacrificing safety goals (Marais, Saleh, & Leveson, 2007). Conversely, in the aftermath of a disaster, the pressures from the media, the public and regulators may increase managerial drive towards the attainment of safety goals, thus relieving pressure for cost-effectiveness (Marais, Saleh, & Leveson, 2007). In this phase, regulators may even impose strong restorative measures to operators, such as suspending operations.

These considerations highlight not only the constant tension for both safety and cost effectiveness that can be found in complex, safety-critical organisations, but also the difficulty in achieving a stable balance between the two over time. For instance, Farjoun (2005) suggested that organisations are normally biased towards cost effectiveness. This is the case for at least two reasons. First, productivity objectives can be easily tracked, as they are usually identified by a set of quantitative productivity-related key performance indicators. Second, the feedback resulting from productivity improvements can be appreciated in a relatively short period of time. In contrast, safety related objectives are not necessarily reflected in quantitative key performance indicators, as many safety issues are qualitative in nature. Furthermore, the time lag necessary to assess the result of safety improvements is usually longer than that required by productivity improvements. As a result, it is easier for an organisation to be biased towards the pursuit of productivity objectives, unless an accident occurs. Farjoun (2005) noted that in these cases there is an increased managerial concern to allocate further resources to attain safety goals, hence leading to a decreased number of near misses and close calls. However, as safety records improve, resources tend to be directed away from safety concerns towards the achievement of productivity goals. Consequently the organisation risks regressing to a state of increased vulnerability, hence allowing another disaster to occur and repeating the cycle. Farjoun's ideas received support from the model of Marais et al. (Marais, Saleh, & Leveson, 2007), who, in an analogous manner, argued that organisations cycle systematically through phases of high efficiency and low thoroughness and vice-versa.

It remains to be investigated whether similar ideas apply to the management of HAI issues in safety-critical domains. For instance, how does the repeated cycling of organisations between states of heightened safety and heightened efficiency affect the management of safety-critical automation? And, in particular, how does such a cycle affect the effective responses to the accompanying HAI issues? Addressing questions such as these requires

examining the problem of HAI issues from an OS perspective.

### 3.2.2. Induced by decisions located at different levels of society at different points in time

One idea implicit in the previous discussion is that the potential for organisational failure and success is actually a collective phenomenon: it reflects the aggregated result of multiple organisational decisions scattered across time and space. This link between (biased) organisational decisions at the blunt end and disaster has been effectively illustrated by Rasmussen and Svedung (2000) in relation to the disaster of the Herald of Free Enterprise, a ferryboat that capsized moments after leaving the Belgian harbour of Zeebrugge in March 1987, killing 193. In relation to this accident, Rasmussen and Svedung highlighted the network of multiple biased organisational decisions that resulted in the capsizing (Rasmussen & Svedung, 2000). These decisions were scattered across various decision makers, belonging to separate departments, from harbour design to cargo management and vessel operations, and were made at different points in times independently of each other (see Figure 1). While each of these decisions was made under stress to meet the “local” demands for cost effectiveness, their aggregated effect gradually pushed the work system—ferry boat operations—into a heightened state of risk.




Figure 1. The pattern in organisational decision making leading to the Zeebrugge accident (Rasmussen & Svedung, 2000).

### Chapter 3. The organisational safety perspective

Rasmussen and Svedung's explanation of the Zebrugge accident is rooted into an expanded view of accident causation that embraces the entire socio-technical control system overseeing safety in a given domain (Cook & Rasmussen, 2005; Rasmussen & Svedung, 2000; Rasmussen, 1997). To Rasmussen, risk management is a societal control problem, in which the control system is composed of the ensemble of actors, located at different nested levels of society, whose decisions, made at different points in space and time, and independent of each other, might ultimately influence safety at the sharp end. As depicted in Figure 2, commencing at the bottom, these actors are usually found at the levels of operational staff, company management, regulators (national and supra-national) and associations, and government. Each of these levels needs to adapt to its own environmental and competitive pressures—such as the changing political climate and public awareness, changing market conditions and financial pressures, changing competency levels—which are normally found in a competitive and dynamic society (Rasmussen & Svedung, 2000). While “low risk operations depend on proper coordination of decision making at all levels” (Rasmussen & Svedung, 2000, p. 11), different actors might lose sight of how these local adaptations influence system safety.

In particular, the model suggests that the higher in the hierarchy a given actor is located, the lower her/his proximity to the operational work system, and the lower the awareness of the consequences of her/his decisions on the work system. For people other than front-end operators, it is very difficult to be fully cognizant of how their decisions eventually affect operations. The farther the distance from the operational level, the more expensive and time consuming it is to gather updated knowledge regarding this level.

These ideas find support in the ethnographically-informed model of healthcare IT proposed by Balka and Kahnamoui (2004). In particular, the two authors averred that the management of a healthcare information system should not be understood by focusing on the level of design alone; there are, in fact, other areas of organisational activities—“social arenas”—which influence and constrain the degree of local intervention that designers can have on any given system. To Balka and Kahnamoui (2004), social arenas include the organisational levels of regulations, politics, international standard development bodies, manufacturers, and national health agencies. These levels, although distant in time and space from operations, may actually influence technology usage. Because of this enlarged perspective, Balka and Kahnamoui seem to extend the range of applicability of Rasmussen and Svedung's model to also include the management of safety-critical technology; warning, in fact, that the understanding of the management of HAI problems requires understanding higher-level organisational activities.






Figure 2. The socio-technical control structure engaged in the control of risk (Source: Rasmussen & Svedung, 2000).

What Balka and Kahnemoui’s model seems to miss in comparison to Rasmussen and Svedung’s model is a commitment to system theory. Such a commitment is certainly retained and further emphasised in Leveson’s Systemic-Theoretic Accident Model and Processes (STAMP) model (2012). Grounded in Rasmussen’s socio-technical model of risk, STAMP provides a method to model complex socio-technical systems as controlled by a hierarchical control structure, a structure characterised by multiple organisational actors at different hierarchical levels, as well as controls, and of feedback loops across these levels. The model characterises failure as the inability of organisational actors at higher hierarchical levels in the hierarchical control structure (e.g., government and regulators) to enforce adequate safety constraints at lower-levels (e.g., operators and manufacturers). The hierarchical control structure involved in safety management is not fixed, but varies across safety domains; thus, STAMP provides a means to model and characterise this. Examples of STAMP have been applied to model the hierarchical control structure involved in the management of the US Air Traffic Control System, and of military

and space operations.

An important aspect emphasised by Rasmussen's, Balka and Kahnamoui's, and Leveson's models is that the control of risk is not the domain of individual organisations alone; rather, it is a societal problem concerning interactions across different hierarchical levels and boundaries in a socio-technical system. This expansive move challenges not only the scope of the current perspectives over HAI issues (reviewed in the previous chapter), but also the scope of many models of organisational safety; as bemoaned by Schulman et al., these tend to focus on single organisations, hence missing higher-level organisational dynamics (Schulman, Roe, van Eeten, & Bruijne, 2004).

Most importantly, Rasmussen and Svedung (2000) suggested avoiding the study of any given socio-technical system from a compartmentalized perspective. They averred that the different levels of the hierarchical control structure have been the objects of study of independent research disciplines. Commencing with the highest level, the government and regulatory activities have been the object of study of political science, economics and sociology; company behaviour has been the object of study of organisational sociology, economics and decision theory; management has been studied by industrial engineering, management and organisation theory; and the level of the sharp end operators has been studied by psychology, human factors, and human computer interaction. Finally, the lowest level, that of the operations, has been the object of study of mechanical, chemical and electrical engineering. However, such a compartmentalised theoretical landscape implies that research disciplines, especially those at the higher-levels, have often omitted the consideration of how dynamics at one level may cascade down and influence safety at lower-levels. For this reason, Rasmussen and Svedung (2000) called for cross-sectional studies into the vertical integration of different hierarchical levels of a hierarchical control structure (Rasmussen & Svedung, 2000). However, they did not accompany this suggestion with any methodological considerations regarding the conduit of such studies; they actually warned that they might require an intense effort, and that specific methodologies do not yet exist. Unfortunately, no work was retrieved from the literature that addressed this gap.

### **3.2.3. An incremental process**

A further important trait of organisational drift into failure is its incremental nature. Organisational drift is not something that happens overnight: accidents are usually preceded by a history of threats and anomalies being downplayed for an extended period of time by organisation members (e.g., Dekker, 2012; Vaughan, 2009; Woods, 2005;

### Chapter 3. The organisational safety perspective

Turner & Pidgeon, 1997).

The idea of drift as an incremental process has its theoretical precursors in the seminal works of Barry Turner and Diane Vaughan. Grounded in the analysis of 87 industrial accidents that occurred in the UK, Turner's *Man Made Disaster* theory (Turner & Pidgeon, 1997; Turner, 1976) viewed accidents as preceded by a period during which the organisation—under existing (biased) set of beliefs—downplays emerging risks and near misses. This *incubation period* may last several years and is characterised by the systematic misinterpretation and disregarding of apparently unrelated hazardous events. Hence, such events accumulate in the work system unnoticed by organisation members; these events do not become a safety target, nor stimulate a revision of existing organisational models of hazards. In this phase, the organisation develops its *failure of foresight*: while resting on a myopic view of the world inconsistent with the way the world really is, it is unable to foresee the incoming failure.

By introducing such an interpretation of disaster, Turner was arguably the first social scientist to show that accidents are not just a physical phenomenon in which an uncontrolled release of energy is involved (such an explosion, a derailment, a collision between two vehicles). Rather, accidents are preceded by a history of hazards being downplayed by the organisation, i.e., they follow the development of a breach, or cultural laceration, between the frames or worldviews maintained within by organisation members and the reality of the world (Pidgeon & O'Leary, 2000). The organisation becomes myopic to signals of danger, and it is only in retrospect that the significance of those signals will be fully appreciated. Therefore, for Turner, the remedy against such situation is the development of a strong safety culture.

The view of failure as an incremental process finds support in Reason's model of resident pathogens (Reason, 1990, 1997). To Reason, the potential for failure in complex organisations follows a pattern similar to the accumulation of pathogens in the human body. Such pathogens may accumulate for years without necessarily causing sickness; it is only when they combine with local triggering factors, such as stress, exposure to toxic chemicals, that they may lead to disease. Similarly, faulty conditions in hardware, software, procedures, and the work environment—i.e., conditions which predispose the organisation to catastrophic failure—may accumulate unnoticed in the work system for years. It is only when such conditions combine with local conditions that they trigger the development of accidents and disasters. One of the merits of this vision is that it stresses the fact that sharp end operators often inherit the outcomes of bad decisions made at the

### Chapter 3. The organisational safety perspective

blunt end, decisions which may increase the potential for errors and mistakes. Therefore, from this perspective HAI issues themselves can be seen as a form of resident pathogens, or *latent conditions* (according to Reason's terminology), in that their occurrence does not necessarily lead to an accident; however, when combined under specific conditions, e.g., high temporal demands or high uncertainty, they may induce human operators to commit fatal errors or make the recovery from emergency situations more difficult. One limitation of this view, however, is that while characterising the role of latent conditions, such as HAI issues, in accident causations, it does not provide guidance on how to investigate the organisational precursors that instilled such conditions in the first place.

Further support for Turner's man-made disaster theory comes from Diane Vaughan's seminal investigation into the NASA Challenger and Columbia space shuttle disasters. An important explanatory concept formulated by Vaughan is the *normalisation of deviance* (Vaughan 2009, 2004, 1997). This phenomenon refers to the fact that, under some circumstances, it becomes routine for engineers and managers to take for granted technical conditions that deviate from expected design performances, instead of viewing them as warning signals (Vaughan, 2005, 1997). In the case of the Challenger, for instance, the anomaly of the O-ring erosion had a history dating back to many years prior to the occurrence of the accident. During this period, the anomaly was reinterpreted as an acceptable and non-deviant condition in the engineering analysis conducted prior to the launch (Vaughan, 2005). Vaughan found the very same organisational pattern repeating in the case of the Columbia. In both accidents, neither of the two respective anomalies became manifest in a clear or immediately understandable manner prior to the disaster. Instead, engineers were exposed to (i) mixed signals, i.e., signals indicating a potential danger that were followed by either less or no damage, reinforcing the belief that the system was safe to fly; (ii) weak signals, i.e., signals regarding risks that after analysis were deemed so unlikely that there was very low probability for them to recur; and (iii) routine signals, i.e., signals relating to events that, while being dangerous, recur routinely with no accident happening, thus reinforcing the idea that the system was operating as predicted. Ultimately, these dynamics contributed to generate the cultural belief that the system was operational and was safe to fly, hence leading the organisation to accept more risk than it realised.

These considerations call into question the collective interpretations and views that guide organisational responses to HAI issues in safety-critical organisations. HAI issues may not be immediately clear to stakeholders at the blunt end of safety-critical organisations, at least not as clear as for those sharp end operators who use the technology on a daily basis.

### Chapter 3. The organisational safety perspective

Vaughan also deserves credit for shedding light onto some macro-dynamics inductive to the normalisation of deviance, namely a *culture of production* and *structural secrecy*. The former reflects cultural and decision-making processes not necessarily internal to the organisation. Both the Challenger and Columbia accidents reflect the broader NASA's transition from (i) a culture of technology to (ii) a culture of production (Vaughan, 2005). The former marked the Apollo era, during which NASA was accustomed to receive a yearly blank cheque from Congress. This promoted a strong reliance on internal expertise and technical positivism. On the contrary, the culture of production was promoted by the fact that the annual budget for the Shuttle Program was allocated under the assumption that the shuttle was an operational vehicle and able to operate on a regular basis, therefore generating commercial return and repaying its costs. The pressure arising from such political accountability was reflected in an increasing focus on meeting schedules and deadlines, which prevented a thorough hazard analysis (Vaughan, 2005). Hence, the culture of production favoured the continuity of launches rather than halting or delaying operations.

Also, the culture of production led to the external contracting of most technical activities. The result was a complex NASA/contractor system composed of several actors, transactions, and different technical languages. This required a burgeoning bureaucratic structure for the purposes of oversight (Vaughan, 1990; 1997), which increased the emphasis on the culture of adherence to hierarchy, procedures and protocols by NASA personnel, in turn reinforcing the belief that the shuttle was safe to fly simply because personnel had followed the rules.

The idea of *structural secrecy* captures the fact that it is very hard for regulators to understand the internal operations of controlled organisations. Structural secrecy arises from factors such as the fragmentation of information, presence of organisational boundaries, limited access to relevant sites and resources, impossibility to observe behaviours or replicate independent testing. In relation to the two NASA disasters of Challenger and Columbia, Vaughan noted that structural secrecy obscured problem seriousness from both top administrators and regulators, so that the belief that it was safe to fly with the O-ring erosion and foam debris prevailed until the occurrence of the two tragedies.

From a methodological perspective, it is noteworthy that Vaughan was able to draw on extensive (i) micro data regarding key decisions taken along the historical trajectory of an anomaly in the phases preceding the launch and the underlying mind-set as well as (ii) macro data regarding the broader institutional context, and to then link these two levels. This was a progression in the field forward compared to Turner, who did not have macro level data as he had relied mainly on official accident reports (Vaughan, 1997). Ultimately,

Vaughan's resulting explanations of both NASA disasters is structural, in that it explains actions of the past as constructed, formed, or organized by major institutional forces, rather than by unconstrained individual choices (Parsons, 2007).

Most importantly for the present thesis, both Vaughan and Turner were effective in showing that (i) the potential for organisational drift into failure is built several years in advance with respect to the final accident; and (ii) that in this process organisations tend to accept, incrementally, increasing deviations from the original norm, the basis for this being that "[e]ach step away from the original norm that meets with empirical success [due to the absence of obvious signs of the violation of safety] is used as the next basis from which to depart just that little bit more." (Dekker, p.6). These properties of organisational failure have been further corroborated by Snook's investigation into the accidental shooting down of two U.S. Army Black Hawk Helicopters over Northern Iraq (Snook, 2002). Snook showed how the friendly fire was actually triggered by an accident-prone context—an organisational and operational context resulting from a stream of different modifications made to the work system at different points in time. While each modification appeared sound according to the local rationality of the actor who made it, their aggregated detrimental effect became evident only when the accident occurred.

It remains to be explored if and how decisions, which might appear reasonable according to the specific situational constraints, when combined might indeed lead to the introduction or poor mitigation of undesirable side effects of safety-critical automation.

### **3.2.4. An interpretive phenomenon**

One further merit in both Turner's and Vaughan's models is that of having highlighted the interpretive nature of organisational drift into failure. Anomalies may become manifest as ill-structured and poorly-defined problems, which are open to several interpretations by organisation members; interpretations that are regulated by the institutionally-accepted interpretive frames (Morgan, 2006; Edmondson et al., 2005; Miliken et al., 2005; Weick, 2000; Turner & Pidgeon 1997). Such frames, or worldviews, regulate sense-making and causal attributions in organisations. They are composed of the set of assumptions, norms, and accepted rules (shared among colleagues) that distinguish what is acceptable and rationale within the organisation from what is not.

These views have at least two major implications for the understanding of organisational failure. First, existing frames or sets of premises orient people towards which signals of dangers to consider, but also towards which signals to ignore, thus leaving potential for the construction of joint blind spots (Hutter, 2005). Therefore, one important implication is that improving safety requires organisations to reason within the accepted frames of

reference, but also to step out of these, realising that blind spots exist, i.e., that certain hazards are not sensed under existing potentially obsolete assumptions, and to then update these frames in an effort to remain sensitive to the likelihood of failure (Woods, 2005b; Woods & Hollnagel, 2006; Pidgeon & O'Leary, 2000).

Second, a lack of consensus about which appropriate frame to activate might develop when facing ambiguous issues. In these cases organisational decision makers need to select an interpretive frame from several (possibly) conflicting frames, in ways that may not necessarily be functional to safety. For instance Milliken et al. (2005) noted that in the years preceding the loss of the Columbia, management rejected repeated requests by engineers for additional images of lift-off, under the conviction that the risk of foam strikes could be minimised considering successful past flight history. In contrast, engineers advocated the need to acknowledge uncertainty and gather further data regarding the likelihood of foam strike. For both groups there was a signal of inadequacy in the established ways to deal with the problem (Milliken et al., 2005); for managers, it was the resurfacing of the request for additional imagery, for engineers, it was the repeated refusal to disclose the required images. Ultimately, since management did not believe a problem existed, the issue remained dormant in the system until the disaster occurred.

To Milliken et al. (2005), this case indicated that powerful elites in organisations might resolve any lack of consensus regarding alternative interpretive views of a possible threat by imposing their dominant frames. In fact, informal and formal power relationships regulate sense-making efforts in organisations, by either blocking or favouring the advancement of selected information and perspectives (Milliken et al., 2005). Milliken et al. (2005) observed that these dynamics cannot be avoided; however, structured debate methodologies (e.g. Schweiger et al., 1989) can be adopted to make the various parties' assumptions and viewpoints explicit, in order to take more informative safety related decisions.

### **3.2.5. Induced by the intrinsic complexity of complex, safety-critical systems**

The potential for organisational failure can also be traced to the side effects related to the introduction of safety defences in support of safety, such as physical barriers, procedures, and alarms. These efforts tend to make systems more complex, and less open to external scrutiny (Dekker, 2011, p. 111). The more systems grow in size, in the number of functions they serve, and in the dependencies they establish with other systems, the higher the number of incomprehensible and unexpected interactions they will experience (Perrow,

1999, p. 72). Such interactions are non-linear and not immediately understandable by personnel and managers. Perrow notes, however, that such *interactive complexity* is not per se a sufficient system property to make catastrophic failure possible. Many organisations exhibit a high degree of interactive complexity, such as hospitals and universities, and also in these contexts many interactions that occur among human operators that are not directly observable. However, in these contexts, a failed activity does not lead irremediably to failure. To Perrow, it is the coupling of interactive complexity with a second system property, i.e., *tight coupling*, which creates the potential for catastrophic failure. Tight coupling relates to the degree of formality of system interdependence. It occurs when production sequences are rigidly defined and cannot be altered. Tight coupling is frequently found in safety-critical systems such as nuclear power plants, in aviation, and in air traffic control. Here, tight coupling creates conditions that allow small failures to propagate quickly across system components and escalate to an accident in unintended and unknown ways. Notably, these insights lie at the heart of Perrow's *Normal Accident Theory* (NAT), a theory that views disaster as a normal condition of modern organisations, due to their tight coupling and interactive complexity.

Although not designed for understanding HAI issues, NAT has been adopted by Tamuz and Harrison (2006) for analysing computerised physician order entry (CPOE) systems, normally used in the healthcare domain. In Tamuz and Harrison, NAT is effective in highlighting the increased coupling of the medication ordering process that resulted from the introduction of the CPOE. This creates the potential for infrequent but potentially harmful and fast travelling errors. For instance, flawed decision rules programmed into the CPOE may quickly affect several patients simultaneously. Based on these considerations, Tamuz and Harrison concluded that organisations should be aware of the trade-offs involved in making their practices more tightly coupled. One limitation that transpires from these considerations is the fact that NAT, although motivated by the intent to avoid to reduce explanations of failure to failures and human error occurring at the system sharp end, is essentially oriented towards the analysis of the operational work system: the dynamics that it characterises are those inherent in the systemic interactions between the technologies and people that can be found at the sharp end of the system, not at the blunt end. Such interactions have a broader and more distributed scope than the HCI models reviewed in the previous chapter (§ 2.2), as NAT does not focus on the unit of analysis as composed by the user and its technology. It is in fact able to consider the propagation of the side effect of technology across a work context. However, the theory does not consider the organisational precursors to failure such as pressures for



cost-effectiveness, poor safety practices, and lack of management support (Hopkins, 2014).

Nevertheless, it must be noted that one merit of the NAT is that of highlighting an important paradox in the management of complex safety-critical systems: a highly interactive system demands a decentralised form of control, as standardised practices cannot cover all of the non-routine situations faced by practitioners (ad hoc solutions need to be devised at the local level of authority). In contrast, effective tight coupling requires a centralised form of control, in order to maximise coordination among system components. One problem is that organisations, in principle, cannot simultaneously be highly centralised and decentralised.

### **3.3. AVOIDING ORGANISATIONAL DRIFT INTO FAILURE**

---

The theories discussed in the previous sections, the NAT in particular, may present organisational disasters as unavoidable, making one wonder why organisations do not experience catastrophic situations on a regular basis (Busby & Bennett, 2007). Countering this view is a body of literature focused on the study of the so called high reliability organisations (HROs). La Porte and Consolini considered HROs as those organisations for which the cost of failure is greatly disproportional to the value of the lesson learned through experiencing the failure (La Porte & Consolini, 1991). These organisations are continuously exposed to significant safety risks, and cannot reduce these risks simply by decreasing external demand, nor can they act on the external socio-political environment. Thus, their ability to achieve a highly reliable performances reflects properties of their internal structure.

Rather than focusing on explaining failure, HRO scholars address questions such as how are risks actually managed? How do organisations achieve a healthy safety record? HRO researchers have addressed these questions by conducting ethnographic oriented case studies of air carrier (Rochlin, 1989; Rochlin, La Porte, & Roberts, 1987), air traffic control (La Porte, 1988), and nuclear power plant operations (P. R. Schulman, 1993). These cases showed that the ability to withstand organisational drift lies in a set of organisational practices which were exhibited by the studied organisations, such as:

- *Top management commitment to safety.* This condition essentially concerns the sensitivity to safety objectives exhibited at the higher ranks of the organisation. If top management is sensitive to safety, it is far easier that lower-level staff will view safety as a priority. Furthermore, top managers sensitive to safety are able

### Chapter 3. The organisational safety perspective

to sacrifice productivity or efficiency goals in order to promote safety goals. This is not necessarily the case when top managers lack safety sensitivity and do not see the safety implications of their decisions;

- *Redundancy.* Another characteristic of high reliability organisations is the redundancy of key operations. Usually achieved by means of the cross-checking of important decisions and redundant communication channels, redundancy ensures that relevant safety-critical decisions are better scrutinised, and that backup systems are in place in case of failure;
- *Adaptable organisational structures.* High-reliability organisations usually enjoy an adaptable organisational structure. In other words, these organisations can maintain a functional organisational structure, i.e., a classic vertical command and control structure in which roles and responsibilities are clearly defined—during routine or normal situations. However, in the case of emergencies, these organisations can adopt a more flexible, decentralised structure—one in which decision making is delegated to the experts who are closer to the safety risk that is being managed, regardless of their rank in the functional structure;
- *Organisational learning.* Organisational learning is another practice that is promoted in high-reliability organisations. This is interpreted as developing safety knowledge in an incremental and controlled way. This implies, for instance, investing in simulations, imaginative exercises, and the study of minor failures in order to envision potentially larger, more severe, ones.

In addition to these developments, recent HRO work has emphasised the specific organisational routines that allow organisations to effectively anticipate the potential for future failures. High reliability organisations have, in fact, the ability to continually challenge their beliefs about the safety of the organisation (Weick & Sutcliffe, 2011). In particular, they do this by:

- *Maintaining a constant preoccupation with failure,* i.e., by paying attention to small signs of danger, and by assuming that such signs may actually conceal larger patterns of failure;
- *Being reluctant to simplify.* Reluctance to simplify refers to the tendency to dismiss simplistic explanations of failure, such as reducing explanations of failure to the occurrence of a human error. Instead, high-reliability organisations tend to favour explanations that consider the actual and specific task demand and contextual factors that lead an operator to commit an unsafe act. One important aspect of the reluctance to simplify is the intent to search for multiple viewpoints,

on the basis that the more viewpoints considered when facing ambiguous issues, the more thorough internal explanations of issues and risks will be;

- *Maintaining a deference to expertise.* This implies ensuring that managerial decision making is not informed by administrative and managerial concerns only, but receives adequate expert input. In other words deference to expertise implies that the expert views of those closer to day to day operations are effectively channelled at the higher-levels of the organisation.

Compared to the work discussed in the previous section (§ 3.2), HRO scholars have been effective in highlighting the sources of positive safety performances, thus suggesting that some effective safety-enhancing routines can be enacted within organisations in order to avoid disaster. It remains to be investigated whether similar routines may promote the effective mitigation and control of HAI issues accompanying the automated tools operated by these organisations.

### **3.4. DISCUSSION: OUTLINING THE RESEARCH GAP AND OBJECTIVE**

---

This chapter has reviewed foundational work in the OS area, and has shown the kinds of insights that the approach can project over the organisational sources of failure and successful safety management. OS's theories provide a rich view of organisational potential for failure and disaster. These theories suggest that organisational failure can be characterised as:

- A by-product of normal administrative processes, i.e., something that results from organisations doing their normal work, not something exceptional;
- Induced by decisions at different hierarchical levels of the organisation;
- An incremental process;
- An interpretive phenomenon; and
- The result of the intrinsic complexity of safety-critical systems.

Additionally, and in contrast, HRO scholars have examined the organisational routines that characterise the organisation ability to maintain successful safety performances—i.e., to avoid processes of drift into failure. These include dynamics such as top management commitment to safety, redundancy, and adaptable organisation structure, as well as specific routines that allow organisation members to continuously challenge their beliefs in the assumed safety of organisational performances.


Overall, the OS perspective would appear to be a valuable approach to adopt to study the organisational precursors to HAI issues. Compared to the traditional approaches reviewed

in the previous chapter (in § 2.2 and § 2.3 respectively), the OS perspective draws its unit of analysis around the enlarged unit of the organisation (rather than the man machine unit and the automation system lifecycle). Thus, if applied to HAI issues, this perspective promotes the followings:

- (i) **HAI issues can be seen as the symptoms of deeper organisational (and problematic) conditions**, rather than purely issues in the design or in the lifecycle process of human machine systems. In other words, the OS perspective suggests to avoid to reduce explanation of HAI issues to problems found in the components of the human machine unit, and in problems in the implementation of specific user centred, human factors and safety assurance principles and practices along the system development lifecycle;
- (ii) **HAI issues can be the starting point for investigations into deeper organisational dynamics**. Assuming that HAI issues may reflect broader organisational issues calls for shedding light into such conditions. If such organisational precursors can be identified, they could provide a wider scope to safety managers, regulators, and policy makers for the mitigation of HAI issues—a wider scope for improvement able to address the organisational roots of the problem and complementary to the scope identified by the human machine system and lifecycle process perspectives.

However, some limitations can be observed. First, the OS perspective has grown in order to allow the understanding of accidents and disasters, but not the problems of HAI issues specifically. Thus, its value has not been demonstrated in relation to the understanding of HAI issues. Furthermore, demonstrating this value is not an easy endeavour. Different models from the OS perspective not only identify different types of organisational precursors, but are also rooted in different conceptual metaphors. For instance, Reason's model is rooted in a view of the organisation and failure as the human immune system; Turner's man-made disaster theory, in a view of the organisation as a cognitive system; Rasmussen's and Leveson's models, in a system theory view of the organisation; Perrow's NAT, in a complexity theory view of the organisation. While each of these theories captures some important facets of the organisational potential for failure and success, there appears to be no single overarching dominant theory that provides an all-purpose perspective (regarding accidents and safety) that can be directly applied to the study of HAI issues. None of the reviewed theories seems to capture the entirety of the intricacies of organisational drift; thus, drift into failure remains a useful conceptual metaphor grouping several theories of failure, each potentially limited in some respects with regard to their ability to capture the phenomenon of interest.

These considerations indicate a theoretical gap that exists at the intersection of the current perspectives, and the OS perspective, as shown by Figure 3. While the former clearly identifies different types of HAI issues and their design/contextual sources (HCI, § 2.2), and their possible process precursors (System lifecycles perspectives, § 2.3), they ignore, however, the broader organisational sources of the problem. In contrast, such sources have been clearly investigated and theorised by the OS perspective—although in relation to accidents and disasters, not to HAI issues per se. Thus, the idea of organisational precursors to HAI issues is situated in between at least three distinct areas of research that do not necessarily appear to relate to each other, so that the theoretical landscape in between lies largely uncharted. This calls for an exploration of such a landscape.



**Figure 3.** Theoretical gap addressed by this research.

Based on the above considerations, the broader aim of this research was to explore the organisational gap existing between the HCI, system lifecycle and OS perspectives in order to contributing to advancing our understanding of how organisations operating in safety-critical domains may manage (or mismanage) HAI issues. The specific research objective was to develop an emerging, theoretical framework of the organisational precursors to HAI issues that can be found in these organisations.

### **3.5. CHAPTER CONCLUSIONS**

---

This chapter has reviewed fundamental works from the OS area. It has concluded that, compared to the HCI and the system lifecycle perspectives reviewed in chapter 2, the OS perspective has, in principle, the potential to address the deeper organisational precursors to HAI issues. The OS perspective is in fact directly concerned with the organisational sources of safety failure and success. It has provided a language and a set of concepts to characterise and discuss organisational precursors involved in the causation of accidents and disasters, and outstanding safety performances. Such precursors are outside the scope of the HCI and the system lifecycle perspectives. However, it should be noted that the value of the OS perspective for understanding HAI issues has not been demonstrated. The area has been developed to explain the precursors to accidents and disasters, in fact, not HAI issues per se.

Thus, this research aimed to explore the theoretical gap existing between the current HCI and system lifecycle perspectives, and the enlarged OS perspective. The objective was to develop a framework of the organisational precursors to HAI issues that can be found in safety-critical domains. The next chapter will describe how this objective was pursued.

# Chapter 4.

## RESEARCH STRATEGY

### 4.1. CHAPTER INTRODUCTION

---

The present chapter describes and justifies the research strategy chosen for developing the the organisational precursors to human automation interaction issues (OPHAI) framework. The chapter encompasses both the overall research plan by which the research objective is achieved, and the lower-level methodological aspects of data collection and analysis.

The chapter is organised as follows:

- Section 4.2 provides an overview of the research strategy;
- Section 4.3 provides the rationale behind the chosen research strategy;
- Sections 4.4, 4.5 and 4.6 describe the type of data, the data collection and data analysis procedures used, and the qualitative validation strategies used within the three studies.

## 4.2. OVERVIEW

The research consists of three studies: two successive retrospective qualitative case studies followed by a third corroboratory subject matter expert (SME) study (as reported in Table 1).

**Table 1.** Overview of the three studies. (Specific details about data, data collection and data analysis procedures are reported in § 4.4 and 4.5.)

	<b>Study 1</b>	<b>Study 2</b>	<b>Study 3</b>
Type of study	A retrospective case study focused on the institutional debate between the National Transportation Safety Board (NTSB) and the Federal Aviation Administration (FAA) regarding the human automation interaction issue (HAI) of the nuisance alerts related to the minimum safe altitude warning (MSAW) system.	A retrospective dual-case study focused on the experience of managing the MSAW-related nuisance alert problem found within two European air navigation service providers (ANSPs).	A subject matter expert (SME) study seeking feedback from safety net experts on the initial OPHAI version delivered by Study 1 and Study 2.
Period referred to by the data analysed	1977-2006	1990-2007	1990 to present
Type of data (more details on § 4.4)	-Documentary data sources	-Interviews, field notes collected during site visits. -Documentary data source	-Qualitative questionnaires filled in by 11 safety net experts
Data Analysis (more details on § 4.5)	Framework analysis (Taxonomic and comparative coding)	Framework analysis (Historical-evidentiary coding)	Framework analysis (Subsuming coding)

The three studies pursued the research objective in an incremental, cumulative way: they all addressed the same research question, each covering aspects not addressed by the other studies. In particular, the first two cases developed an initial version of the OPHAI framework. The third study refined and extended the framework.

This overall research strategy can be considered as a derivative of the case study method



(Yin, 2008). In fact, the three studies have a common context, i.e., the implementation and improvement of an automated alarm from the air traffic management (ATM) domain, the Minimum Safe Altitude Warning (MSAW) system. Developed as a preventive measure against the risk of Controlled Flights into Terrain (CFIT), this alarm has been plagued by one specific HAI issue, namely the frequent generation of nuisance alerts. Thus, the first two studies focus on the organisational trajectory of this HAI issue in the US and the EU respectively. In particular:

- Study 1 is more historical, because it considers the organisational response to the MSAW-related nuisance alert problem between 1977 and 2006. In particular, the case study focuses on the organisational response to the MSAW nuisance alert problem by the National Transportation Safety Board (NTSB), the US agency responsible for the investigation of transportation accidents. The data consist of documentary sources available in the public domain: the set of MSAW-related safety recommendations and safety recommendation letters issued by the NTSB to the US air navigation service provider and regulator, the Federal Aviation Administration (FAA), and the written responses of the latter. This study delivered the first category of the emerging framework of organisational precursors to HAI issues.
- Study 2 is more contemporary. It considers the organisational response to the same problem—the MSAW-related nuisance alerts—as found within two European Air Navigation Service Providers (ANSPs) between 1990 and 2010. This is the case because the MSAW was first introduced in the US in 1977 and then later in Europe in the 1990s. The data for this case study consisted of both data collected during site visits, i.e., semi-structured interviews, field notes, visit reports, internal company documents, as well as data collected from external sources, i.e., interviews with experts external to the organisations, and external documents, however relevant to understand the contexts under study. This study developed a more refined version of the emerging framework of organisational precursors. In particular, in addition to confirming the first category of precursor identified in the previous study, it identified two new precursors.

Study 3 consists of a subject matter expert study (SME) (Marshall, 1996) that was included for the purpose of corroboration. The study sought feedback from a sample of 11 SMEs from the safety net domain. Safety nets are the class of air traffic management (ATM) alarms that the MSAW system belongs to (as it will be further explained in due course). Data were collected by means of a structured group discussion and a qualitative

questionnaire. This study allowed the categories of the framework to be refined and expanded so that it could be delivered in its final form.

Data analysis in all the three studies was based on a common qualitative data analysis approach known as a *framework analysis* (Ritchie & Spencer, 2002). This approach was selected because it provides a high-level data analysis framework that, while specifying basic data analysis steps to be followed in qualitative research, it is flexible enough to be adapted to different research contexts. This flexibility is evidenced by the fact that the three studies used different coding strategies depending on the specific analytical challenge they posed (details about both data collection and data analysis are provided by § 4.4 and 4.5)

While Study 1 relied on data available in the public domain, Studies 2 and 3 were possible thanks to the organisational access provided by the sponsor of this research: EUROCONTROL, the European Agency for the Safety of Air Navigation. EUROCONTROL funded and hosted the first three years of this PhD at its Experimental Centre in Brétigny-sur-Orge, in southern Paris. In particular, years two and three of the PhD were carried out within the Safety Team of the agency. In addition to providing intellectual support for the work, the team was essential in gaining organisational access and supporting site visits in Study 2; and in providing access to the network of safety net experts, SPIN, which was relevant for studies 2 and 3 (as it will be explained further in due course).

### 4.3. RATIONALE

---

This section describes the rationale behind the research strategy outlined in the previous section.

#### 4.3.1. Rationale for grounding the research on the case study approach

The research strategy adopted in this research can be considered as a derivative of the case study approach. The case study approach is usually defined as a research strategy that allows a researcher to concentrate on a specific situation bounded in space and time, because it contains some elements worth discovering (Yin, 2008). During the investigation, these elements are closely explored in their natural setting, in the absence of control conditions, in order to provide a detailed multidimensional picture of the analysed situation. The case study approach has been applied extensively across a range of subject areas such as business and management, information system, innovation management, history and anthropology (Dubois & Gadde, 2014; Yin, 2004; Benbasat, Goldstein, & Mead, 1987).

Most importantly, the case study is an accepted approach for investigating processes of organisational drift. It can be found, in fact, in the work of many scholars from the organisational safety area (e.g., Roe & Schulman, 2008; Perin, 2006; Snook, 2002; Vaughan, 1997; Shrivastava, 1994; Weick, 1993; Roberts, 1990; Rochlin, La Porte, & Roberts, 1987)<sup>2</sup>. This tendency is understandable considering that these authors were usually driven by a quest for sense making and insight into a real life phenomenon—organisational sources of failure and/or high reliability performances—that encompasses important contextual conditions that cannot be studied at distance, as they cannot be easily isolated from their context. This type of investigation requires an approach that is able to make sense of the multiple operational links and dependencies, and the cultural and structural features that pertain to a specific and complex organisational context. In fact, the emphasis of this kind of investigation is on the discovery and identification of the relevant conditions and dynamics involved in the specific problem situation, rather than on measuring frequencies and establishing generalisability to a wider population.

The case study approach meets these demands because it provides a window into a portion of organisational life bounded in time and space that can be examined in great detail. By using this approach, the researcher can collect many profound insights about insiders' interpretations, issues, relationships and processes at work which are related to the phenomenon of interest (Dalcher, 2004, 2009). The researcher can leverage on multiple data sources, ranging from archival records and internal company documents, to interviews with insiders and records of observations of those people directly involved in the problem situation (Yin, 2009). The data used can be both quantitative and qualitative. Ultimately, such flexibility in data collection usually results in a rich evidentiary base, which in turn maximises opportunities for insight and discovery.

On these grounds, the case study approach was determined to be appropriate for the present inquiry into the organisational precursors to HAI issues. The flexibility of the approach and its focus on discovery were considered to be important properties that cohere with the exploratory purpose of the research.

Note that a valid alternative to the case study approach is the Grounded Theory methodology (GTM) (Glaser & Strauss, 2009). This is a classic qualitative approach, which

---

<sup>2</sup> Appendix A provides an extensive review of the methodological approaches used by these scholars. Despite the variations that can be found across different approaches, they can all be classified as case studies, because they focus on the in-depth analysis of one organisational context. The only exception is provided by Turner, described later in this section, who made use of GTM.

Barry Turner (1981, 1983) brought into the OS area in order to develop its foundational man-made disaster theory (§ 3.2.3). However, GTM is extremely data-demanding, as it requires individual organisational cases to be available for comparative purposes. This is evident in the seminal work of Turner, who in applying the method could rely on the availability of a relatively large sample (>80) of official accident reports published by the British government between 1965 and 1975<sup>3</sup>. Replicating the same approach in the present context was not feasible because similar kind of reports are usually not produced in relation to failed automation programmes or implementations. Unlike major transportation accidents, these failures are not the subject of official investigations by official independent bodies. For this reason the case study method was chosen over the GTM.

### 4.3.2. Boundaries of the research

In this research the boundaries have been drawn around the organisational trajectory (or history) of a selected HAI issue, the *nuisance alert problem*, related to an automated system from the Air Traffic Management (ATM) domain, the MSAW.

This decision was motivated by the fact that processes of organisational drift develop over periods of years, or, in some cases decades, as illustrated by the case of Challenger and Columbia (§ 3.2.3). For this reason, the boundary of the case had to account for such a time period.

In addition, the focus on the organisational trajectory of a selected HAI issue was inspired by Vaughan's seminal investigations into the Challenger and the Columbia accidents. Both investigations are based on the central idea that the way an anomaly is defined, negotiated, and controlled is a reflection of the broader organisational context in which the anomaly is addressed. Therefore, by tracing the history of the organisational events and decisions related to an anomaly one can shed light on the broader organisational conditions in which the anomaly was "processed". For instance, by looking at the history of the events and organisational decisions related to the O-ring erosion anomaly prior to the Challenger accident, Vaughan was able to identify the higher organisational and regulatory conditions that allowed the anomaly to become normalised despite its safety relevance, which in turn led to it remaining unaddressed in the operational system for years before the disaster unfolded (Vaughan, 2004) (§ 3.2.3). On this basis, it was

---

<sup>3</sup> More details regarding Turner's application of the GTM methodology are available in Appendix A.

determined that the same analytical focus could be replicated in the analysis of the organisational trajectory of a selected HAI issue—the nuisance alert problem related to the MSAW system.

### **4.3.3. Selected application case: the MSAW**

The MSAW is a subsystem of the main radar system available within air traffic control centres to alert air traffic controllers of an aircraft's close proximity to terrain. The system compares the current or projected aircraft altitude against a predefined terrain database. Whenever an aircraft descends or is about to descend below a predefined minimum altitude, the system generates a visual and aural warning. In Europe the controller has to inform the pilot of the imminent danger upon reception of the MSAW alert, so that the pilot can take corrective action if necessary. The MSAW is the ground equivalent of the perhaps more famous Ground Proximity Warning System (GPWS), which is available in the cockpit and is intended to warn pilots of dangerous proximity to terrain.

The MSAW is intended as a protection against Controlled Flight Into Terrain (CFIT) accidents, i.e., accidents occurring whenever an aircraft is flown into terrain, obstacles or water without any technical failure and with the crew being unaware of the imminent collision (Wiener, 1977). These accidents make up one of the leading categories of civil aviation accidents (Sumwalt, 2014). The likelihood of CFIT is considered to be higher during the landing and take-off phases of the flight, especially when in the presence of high terrain, and/or man-made obstacles, such as buildings or antennae.

The MSAW was first developed and introduced in the 1970s in the US. In Europe, the tool has been introduced in the 1990s. Here, the MSAW is considered a member of a class of ground-based alarms called *safety nets*, which use primarily surveillance data, and provide warning times to controllers of up to two minutes (SKYbrary, 2014). In addition to CFIT, safety nets are intended to warn controllers of other potential risks, such as collisions with other aircraft (Short Term Conflict Alert or STCA), infringements of protected airspace volumes (Area Proximity Warning System or APW), and deviations from the final approach path (Approach Path Monitor or APM) (EUROCONTROL, No date; SKYbrary, 2014).

Despite several decades of operational use, some aspects of MSAW operation have remained a concern (Howell, 2011). The most notorious problem with the use of the MSAW is its tendency to generate a high rate of false alarms, usually due to the difficulty of parameterizing the MSAW terrain database. In fact, the higher the mismatch between the database and the underlying terrain, the higher the number of nuisance alerts. Technically speaking, such matching requires fine-tuning of the terrain database. This is

more difficult in the presence of high and variable terrains, such as reefs, where the alert can play a more crucial role. Operational evidence has shown that the nuisance alert problem has resulted in the alarm being downplayed or ignored by air traffic controllers even in the presence of real danger; thus contributing, in some cases, to major commercial aviation accidents.

The MSAW system was selected as an application case for the following three reasons:

- *First, it is a good revelatory case.* The MSAW enjoys a relatively long operational history that has been troubled by HAI issues, (the most notorious being the issue of nuisance alerts), and that stretches across nearly three decades in the US and two decades in Europe. As discussed earlier, this temporal scale is large enough to exhibit the dynamics of organisational drift that are of interest here. On the contrary, novel, more recent, automation would not necessarily exhibit a similar history.
- *Second, it is an interesting case for the sponsor of the research.* At the time the research was conducted, EUROCONTROL was in the process of developing guidance material and specifications for safety nets. Therefore, the Safety Team considered the research to be an opportunity to gain useful insights into happenings at the organisational level that complement the on-going efforts in the area of safety nets (which primarily address the technical aspects of implementation and improvement).
- *Third, it is a societally relevant case.* The successful management of the MSAW system directly affects the safety of air transport. Therefore, it makes sense to expand the knowledge base related to the management of this system, as this could be directly beneficial to organisations, manufacturers, standardisation bodies, and regulators involved in its development and overseeing.

#### **4.3.4. Validation and generalisability**

The research strategy employed in this study departs from the single case study approach in order to attain greater validity and generalisability for the emerging framework than could be achieved from a single case study. In particular, three strategies were used at the level of the overall research:

1. Replication of the same type of case study design;
2. Methodological triangulation;
3. Consideration of different organisational contexts.

1. *Replication of the same type of case study design.* Replication of the same type of case study design in two separate, sequential cases (Studies 1 and 2) ensured that the methodological insights and lessons of the first case (Study 1) could be carried forward and applied to the second case (Study 2)—which could in turn be executed in a more controlled way. Furthermore, replication was also functional in exploiting synergies between the two studies: the MSAW has been invented in the US and later deployed in Europe; therefore, it was envisaged that looking at the MSAW in both the US and the EU ATM systems would provide a more comprehensive view of the system's historical background. Finally, the use of two case studies safeguarded against the risk of the failure of one case; a tangible risk in an exploratory investigation like the present study. It is in fact acknowledged that although the exploratory mode of research involves “pushing the frontiers of knowledge” (Phillips & Derek, 2010, p. 59), it also does not guarantee that a particular research project will produce the desired outcome.

2. *Triangulation.* Study 3 was executed to achieve *triangulation* at the level of the research design. This study made it possible to cross-check the quality and validity of the initial version of the OPHAI framework by using an alternative, independent approach than initially used in the first two cases. The chosen format of this study, a SME study, consists of getting information from key informants or experts whose opinions carry weight or plausibility in a given domain. Because of their unique knowledge and understanding, relevant experts can provide extensive insights into a specific topic in the form of additional supporting or disconfirming comments and/or examples. The use of this kind of study is adequate to the qualitative, or interpretive orientation, of this research<sup>4</sup>. In this type of research, the focus is usually on studying small numbers of subjects, often single individuals, single groups, or single organisations, so that truth and reliability are especially important in relation to the viewpoints of the people engaged in the context being studied<sup>5</sup>. Therefore, one important way in which one can achieve validity is by

---

<sup>4</sup> The rationale for grounding this work in the interpretive mode of research is provided under Appendix B.

<sup>5</sup> This is not to say that generalisability is neglected in interpretive research. Indeed, generalisability can remain a desirable property also in this mode of research. However, rather than statistical generalisability, a more plausible way for interpretive research to achieve generalisability is to rest on the argumentative logic of the researcher—argumentative logic about the extent to which research results can be applied also to contexts other than the investigated one(s). Specifically to this research, such an argument for generalisability is initiated by the reminder of this section, and then is further consolidated by sections 7.3 and 8.5.

seeking confirmations, refutations, and reformulations of the interpretations that develop along the course of the research (Morgan, 2006, 1997). This corresponds to searching for *interpretive validity*, i.e. ensuring that the research adequately represents the viewpoints, thoughts, intentions and experiences of the study's participants (Johnson, 1994; Maxwell, 1992). Striving for interpretive validity implies protecting the research against the risk of misconstruction, i.e., the risk of failing to understand and/or misrepresent the meaning of events as understood from an insider's perspective. It is in light of these considerations that the use of an SME study was considered appropriate for the present research. (Note that other strategies were used within each study to increase their validity, as it will be explained in § 4.6).

*3. Consideration of different organisational contexts.* The cases studies found in the OS literature are usually based on the consideration of either single negative cases, i.e., accidents (e.g., Snook, 2002; Turner & Pidgeon 1997; Vaughan, 1997; Weick, 1993), or single successful cases, i.e., the HRO (e.g., La Porte, 1996; LaPorte & Consolini, 1991; Roe & Schulman, 2008). Such a single case study approach would have encountered issues regarding the generalisability of its findings, as these would have been relevant only for the individual studied context. The chosen research methodology addresses this concern by maximising the number of organisational contexts included in the research. In particular, the first two case studies explicitly consider a total of three organisational contexts that have been analysed in depth. Two of these contexts (Study 2) addressed a successful and a less successful case. Finally, the third study is informed by the knowledge of different organisational contexts found in European ANSPs, as reflected in the professional backgrounds of the safety net experts involved in the study.

Having covered the aspects of validity and generalisability, this section has completed the description of the strategic aspects of the research and has set the stage for the description of the lower-level, methodological aspects of the research, which will be addressed in the following two sections.

### **4.4. DATA AND DATA COLLECTION**

---

Considering the qualitative focus on the history of the MSAW, it was considered as appropriate data any key documents or insider accounts reporting on activities, decisions, and events influencing the adoption, implementation, management, human factors and safety assurance and use of the MSAW in the studied organisations. In investigations of organisational sources of failure it is normal to rely on a variety of data sources, such as:



official accident reports, internal company documents (e.g., Ocasio, 2005; Turner & Pidgeon, 1997); field notes from ethnographic observations of operators at work, and transcripts from single or group interviews (e.g., Rochlin, 2011, 1996; Roe & Schulman, 2008; Roberts, 1990); a combination of both official documentation and in-depth interviews (e.g., Snook, 2002; Vaughan, 1997); and, also, on secondary data sources (e.g., Weick, 1993). In exploratory research like the present thesis, flexibility in data collection is desirable because it maximises opportunities for discovery—i.e., it maximises opportunities for finding possible organisational precursors to HAI issues wherever they can be found.

Furthermore, the data had to be representative of the viewpoints of both (i) the air traffic controllers at the sharp end, i.e., the users of the technology; and (ii) the other stakeholders at the blunt end, such as supervisors, managers, R&D directors, safety experts, and international and national regulators involved in the management of the MSAW. The former viewpoint provided insights into the actual use of the alarm in context, and the accompanying HAI issues. The latter viewpoint was instrumental to provide insights into the rationales, interests, and frames behind the (normal) decisions and conditions influencing alarm development, adoption, operation, and optimisation.

Overall, these considerations justify the choice of the type of data that have been employed in this research. The data are described in the next sub-sections.

### **4.4.1. Study 1**

The first empirical study of the research builds on the interpretive historical case study approach described in the previous section (§ 4.3.2). The case study focused on tracing the historical trajectory of the nuisance alert problem within the context of operation of the MSAW system's operation in the US, from the introduction of the alarm in 1977 to 2006.

#### **4.4.1.1. Data sources**

The data sources considered for Study 1 included the following documentary sources:

- The set of NTSB safety recommendation letters (n=11) that reported problems identified by the NTSB with MSAW operations. These letters' average length was six pages, with a minimum of two pages, and a maximum of 13 pages. Usually safety recommendation letter cover the following topics: a description of the evidence pointing to a given safety problem (a problem with MSAW operations in the present analysis); a description of the nature of the problem(s), as understood by the agency; and the presentation of the safety recommendations issued to the

FAA to fix the problem(s) identified. In this study, the analysed safety recommendation letters carried a total of 25 safety recommendations addressed to the MSAW. Appendix C provides an example of one of the NTSB safety recommendation letter and its safety recommendations considered in this study.

The NTSB safety recommendation letters were retrieved from the NTSB safety recommendation database (described in the next section); and they covered the period 1977 (the year the alarm entered operational use) to 2006. No safety letters and recommendations were found after 2006. Appendix D provides the complete list of the NTSB safety recommendation letters and NTSB safety recommendations analysed in this study.

- The correspondence exchanged between the NTSB and the FAA that addressed the implementation of the above safety recommendations. This included a total of 111 letters (see Appendix D for the list of the letters considered by the study), which were on average half a page long. These letters were exchanged as, upon reception of an NTSB safety recommendation, it is a standard administrative practice for the FAA to issue a response letter documenting the actions that it will undertake in response to the recommendation. This is usually followed by one or more correspondence exchanges between the two agencies. Such exchanges are usually triggered by NTSB letters that comment on the FAA's response actions. These exchanges are terminated when the NTSB declares the safety recommendation close. This step can have three outcomes: (i) *acceptable*, when the FAA corrective actions satisfy the intent of the initial safety recommendation; (ii) *acceptable alternative action*, when the FAA implements corrective actions that are alternative to those demanded by the NTSB in a way that still satisfies the initial safety recommendation; and *unacceptable action*, when the FAA response does not satisfy the initial NTSB safety recommendation. Further background information about safety recommendations is provided in 5.2.2.

It is important to anticipate that this data set was analysed three times in three different manners. The motivation for this lies in the relative uniqueness of the kind of data analysed. Only one study (Tasca, 1990) which made use of similar data was found, and this study did not use this data for research objectives comparable to those of the present research. Therefore, it was only after the first data analysis stage—data familiarisation—that it became evident how the initial data set could be actually analysed. (Data analysis for Study 1 is described in § 4.5.1.)

#### 4.4.1.2. Data bases accessed

The above data were retrieved from two publicly available on-line databases maintained

by the NTSB and the FAA. More specifically, the NTSB's safety letters and recommendations were retrieved from the '*NTSB Safety Recommendations and Safety Letters database*'. During the study, the database has underwent some improvements. Thus, a first batch of data was downloaded from the older version of the database (available at: <http://www.nts.gov/Recs/letters/letters.htm>), and a second batch was downloaded from the latest version (available at: <http://www.nts.gov/safetyrecs/private/QueryPage.aspx>). The ensuing safety letters exchanged between the NTSB and the FAA were downloaded from the '*NTSB Safety Recommendations to the FAA with FAA responses*' database (<http://www.asias.faa.gov>).

Although the details of incidents and accidents detailed in most safety recommendations sufficed for the purposes of this study, in two cases it was necessary to access the original accident report to gain access to details not reported in the letters. These were downloaded from the '*NTSB Accident Reports*' database (<http://www.nts.gov/investigations/reports.html>).

### 4.4.2. Study 2

The second empirical study of the thesis study also builds on the interpretive historical case study approach described in the previous section (§ 4.3.2). The case study traced the historical trajectory of the MSAW-related nuisance alert problem within two European ANSPs. In this context, the study considers the perspective of the developers, engineers, managers, air traffic controllers and safety experts involved in the implementation of the MSAW and the management of the related nuisance alert problem.

#### 4.4.2.1. Sample of organisations

The selection of the ANSPs involved in this study was made according to the following criteria. Firstly, ANSPs had to have experience with the introduction and operation of the MSAW system. Secondly, an effort was made to include at least a successful, or best practice, case as well as a less successful one. This would have in fact increased opportunities for insights than the analysis of a single organisational context—as carried out in Study 1. These two criteria led to the identification and inclusion in the study of two ANSPs, both henceforth referred to using fictional names:

- ANSP1, here referred to as *Alphasky*, is a large European ANSP that is recognised as a leader of research and development of new automated systems for ATM. Specifically to MSAW, Alphasky was the first European ANSP to introduce the system and is recognised today as one of the best in class MSAW implementers in the industry;

- ANSP2, here referred to as Deltasky, is a small East European ANSP. It first introduced the MSAW system in 2004, and at the time of data collection it was in the process of improving its MSAW and other safety nets.

Table 2 compares the size of the two organisations.

**Table 2.** *Profile of Alphasky and Deltasky*

	Alphasky	Deltasky
No. of employee	>5000	<1000
Traffic movements on peak days	>9000	<3000

#### 4.4.2.2. Organisational access

Among organisational researchers, gaining organisational access for the purpose of conducting field-work is usually acknowledged to be a main hurdle. This is especially true when seeking access to high-risk organisations. As noted by Bourrier (2011) researchers are often discouraged from accessing these organisations due to factors such as the number of guarantees that the researcher is requested to provide; the requirement to be assisted by an internal surveillance team during the visits, which risks altering the setting under investigation, and the time needed to gain access.

Because of these considerations, careful planning was undertaken to obtain access to Alphasky and Deltasky. In particular, access was obtained according to a two-step process: the first step consisted of seeking access to the EUROCONTROL SPIN network. This international network brings together operational, technical and safety experts from different ANSPs and industry that work together with EUROCONTROL in order to enhance safety net implementations in Europe (EUROCONTROL, No date). In particular, at the time of the study the network was working on defining guidance material and specifications for safety nets. Furthermore, the internal EUROCONTROL SPIN coordination team was active in providing assistance to the ANSPs that needed to optimise their safety nets.

In February 2009, the EUROCONTROL Safety Team manager, the industrial mentor of this dissertation, established initial contact with the leaders of the SPIN network. During a dedicated meeting, these individuals were briefed about the objectives of the research. Access to the SPIN network was subsequently granted, and the researcher was invited to participate in subsequent SPIN meetings. In particular the researchers attended three of these meeting in 2009. Attendance of these meetings was instrumental in gaining initial

direct exposure to the safety net domain, and, most importantly, in establishing informal contact with relevant ANSPs' representatives that could be interested in the research.

The subsequent step consisted of approaching the interested representatives in order to request organisational access. Four representatives of four different ANSPs were sent an e-mail that specified (i) the aim and objectives of the research, (ii) the location and duration of the requested visits, (iii) the support needed, (iv) the data collection activities that needed to be carried out, and (v) the background of individuals the research team wished to interview. (This e-mail also specified that the names of people, companies, and facilities would remain strictly confidential.) Access to Alphasky and Deltasky was obtained as a result of this process. The contacts of these organisations were very supportive with regard to the objective of the research. They granted the research team access to their operational and administrative facilities and permission to interview staff involved in MSAW introduction and operation.

### 4.4.2.3. Data collection

Fieldwork took place between February and November 2009. During this period the researcher visited two approach and two tower control centres of Alphasky and one approach and one tower control centre of Deltasky. The visit to the first Alphasky site was conducted together with a senior safety expert from the EUROCONTROL Safety Team, who was instrumental in facilitating access to the organisation, and stimulating discussions with participants. The same expert also facilitated an initial teleconference with a representative of Deltasky, although the visit to this organisation was conducted by the author of this research alone.

Fieldwork was guided by the basic rule of interpretive research of trying "to get inside [the] situation to understand it as far as possible on its own terms" (Morgan, 1997, p. 301). The research team approached the field from a learner's perspective rather than from that of an expert, meaning that they focused on understanding the local experience with the MSAW and the management of nuisance alert problem, as seen from the viewpoint of the study participants. In doing so, an effort was made to leave preconceptions and hypotheses behind and to suspending judgment, in order to avoid premature closure and to create room for new hypotheses to emerge.

Semi-structured interviews were an essential part of the field-work. Participants were prompted based on an initial list of open ended questions, allowing for deviations from this list whenever important issues emerged. The participants were encouraged to discuss in detail aspects that were relevant to the purpose of the research. During the interview process, the participant was invited to do most of the talking, while the researcher listen

carefully and took notes. Appendix D provides the interview guides employed for air traffic controllers and back end roles.

The main concern driving the selection of participants was to identify and talk to people involved in MSAW-related activities within their organisations. The goal of the interviews was to understand (i) the viewpoints of relevant stakeholders at the blunt end who had been involved in the decisions to introduce, manage and improve the system and (ii) the viewpoints of the MSAW-users, i.e., operational controllers. Thus, prior to the site visits, a list of possible interviewees was provided by the organisations, and after each interview participants were asked to suggest names of other people that could be potentially interested in the research.

#### 4.4.2.4. Data sources

The study made use of data sources both internal and external to the organisation. Internal data sources included the following:

- Qualitative interviews. A total of 28 staff members across the three sites were interviewed (see Table 3). These included individuals both in back end and front end roles. Interviews were also conducted with 7 participants external to the two organisations. 5 of these belong to the SPIN network, while other 2, a EUROCONTROL policy maker and an EASA regulator, were external to both the SPIN and the studied organisations.
- Field notes. In addition to conducting interviews, the researcher also carried out direct observations of air traffic controllers at work in their control rooms. This made it possible to identify their monitors and displays; to observe how the MSAW system fits into the actual operational environment; and, in particular, to understand how the system is used and how nuisance alerts may occur during its use.
- Visit Reports. During the visits the research team wrote copious notes of what was said and what happened. Within three days of each site visit, a visit report that summarized the main findings of the visit was prepared, checked internally, and then sent back to the study participants in order to confirm the accuracy of the information gathered.
- Documentation. Internal documentation included service notes, internal safety net requirements and guidance material, and MS Power Point presentations. The internal documentation was supplemented with external documentation such as regulations, EUROCONTROL European safety policies and safety net guidance material, and standards by the International Civil Aviation Organisation.

**Table 3.** Study 1: list of participants interviewed.

Interviewed staff	Case 1: ALPHASKY	Case 2: DELTASKY	EXTERNAL EXPERTS
Blunt end roles	<ul style="list-style-type: none"> <li>- Head of ATC (site 1);</li> <li>- Head of ATC (site 2);</li> <li>- Safety net expert;</li> <li>- MSAW engineer;</li> <li>- Former safety net engineer;</li> <li>- Quality of service specialist;</li> </ul>	<ul style="list-style-type: none"> <li>- Project manager;</li> <li>- 2 IT engineers;</li> <li>- Training Expert;</li> </ul>	<ul style="list-style-type: none"> <li>- 2 EUROCONTROL SPIN leaders;</li> <li>- 1 R&amp;D director;</li> <li>- 1 safety net leader;</li> <li>- 1 safety net expert;</li> <li>- 1 EUROCONTROL policy maker (co-author of the ESARR4 policy);</li> <li>- 1 EASA regulator.</li> </ul>
Sharp end roles	<ul style="list-style-type: none"> <li>- 10 controllers (site 1);</li> <li>- 5 controllers (site 2);</li> <li>- 1 supervisor (site 2).</li> </ul>	- 2 controllers	<i>Not applicable</i>
Total	<b>22</b>	<b>6</b>	<b>7</b>

#### 4.4.3. Study 3

Study 3 consisted of an SME study. In this study, feedback was collected from a group of 11 experts in the domain of safety nets by means of an initial group discussion, and an individual qualitative questionnaire.

##### 4.4.3.1. Recruiting of experts

In research using experts, the background of the individuals involved is a key component in ensuring the study's validity. Thus, the driving concern was to have a sample of experts with first-hand experience in the safety nets domain. The experts were recruited from the EUROCONTROL SPIN network, because, based on their professional experience, these experts contribute periodically to the various initiatives undertaken by the network in the safety net domains. This ensured that the experts not only understood the organisational patterns for which corroboration was sought, but also that they were able to effectively express their thoughts, opinions and perspectives on these patterns.

##### 4.4.3.2. Data Collection

In an SME study, there are many ways to engage with experts. In the present study, the following combination of method was sought: (i) the research was presented during an event attended by safety net experts; then data was collected by means of (ii) a group discussion and (iii) an individual semi-structured questionnaire returned to the researcher via e-mail.

The event in question was a two-day meeting of the SPIN community, which took place in May 2013 at the EUROCONTROL headquarters in Brussels. On that occasion, a one-hour slot, dedicated to this study, was included in the meeting's agenda.

- *Presentation.* During the allocated time-slot, an initial presentation lasting about twenty minutes was given. Here, the meeting participants were (i) briefed about the purpose of the study and the importance of their involvement; (ii) provided with a description of the initial version of the framework (as resulting from Studies 1 and 2); and (iii) instructed about the kind of feedback that the researcher aimed to collect. It was also emphasised that confidentiality would have been preserved—in particular that no names of individuals, companies, sites or airports would have appeared in the research manuscript and the related publications.
- *Group discussion.* The presentation was followed by a 40-minute group discussion, during which the interested participants asked clarificatory questions about the research, and provided their initial reactions to the categories of the initial version of the framework (resulting from studies 1 and 2). This mode of participant involvement is known to establish trust and rapport with study participants, which in turn increases the richness and validity of participant accounts (Meho, 2006).
- *Questionnaire.* During the event, the experts were provided with a semi-structured qualitative questionnaire, the structure of which is described in the following subsection. 11 experts participated actively in the initial group discussion and returned their completed questionnaires via e-mail between May and June 2013 via e-mail. Each returned questionnaire was usually followed by between one to four e-mail exchanges in order to elicit additional information and clarify relevant issues.

### 4.4.3.3. Questionnaire

The qualitative questionnaire used in this study included 10 open-ended items, organised into 5 sections. The first section included items aimed at collecting biographical data. The second, third, and fourth sections formed the core of the questionnaire. Each section presented a definition of one category of organisational precursor. The definition was followed by two questionnaire items, the first inviting the expert to comment on the definition provided, the second asking for recommendations for improvement (relevant to the organisational precursor just commented upon). The last section included two concluding items. Prior to the study, the questionnaire was designed and refined with the support of a supervisor knowledgeable about the ATM domain.



#### 4.4.3.4. Profile of the SME group

The average years of experience in the safety net domain of the group of experts was about 10 years, with a total cumulative experience of roughly 105 years (see Table 4). The group encompassed the perspectives of the different organisational actors that may be involved on a typical safety net implementation and/or improvement (see Figure 3). Figure 4 shows that the group of experts included individuals who had been involved in the implementation of safety nets (n=6), in the evaluation and improvement of existing safety nets implementations (n=5), and in the development and validation of safety net guidance material (n=5). Only one expert had no direct experience in the safety net domain. Regarding the type of safety nets the experts were familiar with, the majority of them were knowledgeable about the MSAW and the short term conflict alert or STCA (Figure 5). Three of them had also been involved in the requirement definition and the implementation of other automated systems in addition to safety nets.

**Table 4.** Study 3: profile of the group of experts participating in the exercise.

N. of participating experts	11
Average experience in the safety net domain ( <i>years</i> )	≈10
Min* (*:excluding the experts with no experience in the safety net domain)	5
Max	14
Total cumulative experience ( <i>years</i> )	105




Figure 4. Study 3: organisational perspectives represented by the group of experts.





Figure 5. Study 3: type of involvement of the participating experts in the safety net domain.



**Figure 6.** Type of safety nets the experts had direct experience with.

## 4.5. DATA ANALYSIS

---

Data analysis used in the three studies was based on framework analysis (Ritchie & Spencer, 2002). Originating in the area of policy research, framework analysis consists of a broad and adaptable high-level framework that can be used for qualitative data analysis across different situations and research problems. The phases of the approach include:

1. **Data familiarisation.** This phase consists of immersing oneself in the (large) body of data that is usually collected in qualitative research. At the end of the familiarisation phase, the researcher should have gained (i) an understanding of the breadth and depth of the available data corpus and data set(s); (ii) an initial and rough feeling about (potential) key ideas expressed in the data; and, most importantly, (iii) an idea about which coding method can be used to analyse the data.
2. **Coding framework development.** The second phase involves the development of one or more coding framework(s). A coding framework consists of a hierarchical set of codes and categories that synthesises the data in a way that is meaningful for the purpose of the research. According to Saldaña (2012), the development of a coding framework involves two related sub-phases, or, more specifically, two levels of coding: first-level and second-level coding. The former involves coding the raw data into first-level categories; the latter, involves coding the first-level categories, so as to deliver higher-level, and more theoretically-oriented categories.
3. **Data display.** The third stage, data display, consists of arranging the data in a way that provides an integrative at a glance view of the emerging themes and/or patterns found in the data. Miles and Huberman (1994) provide a comprehensive review of the kind of data display methods that are useful in this phase of the analysis.

These phases drove the data analysis in the three studies and are described in the following sections. Note that the studies differed from one another regarding the implementation of the core phase of the approach, namely phase 2, *Coding framework development*. This is due to the fact that each study used different data sources that required different coding methods, namely taxonomic coding (Spradley, 1980), comparative coding (Shreier, 2002), historical-evidentiary coding, and subsuming coding (Mayring, 2010). Table 5 (see next page) summarises the different coding frameworks developed in each study and the corresponding coding method(s).

Indeed, coding framework development can be pursued using different coding methods. Coding methods frequently mentioned in the literature include open, axial, and selective coding—the three coding methods most typically associated with GTM (Glaser & Strauss,

2009). However, coding should not be restricted to these approaches. Other methods are available, and Saldaña’s comprehensive coding manual provides a review of more than thirty of them (Saldaña, 2012). He suggests that this variety can be attributed to the fact that coding methods should be fitted to the unique goal and demands of the research at hand. Thus, in the present research, different coding methods have been used depending on each study’s data set(s) and analytical goal.

While each specific method is described in the following section, it is worth noting that Table 5 lists a total of five coding frameworks—and not just three. This is due to the fact that Study 1 delivered a total of three frameworks and not just one: as mentioned during the presentation of the data sources used in this study (§ 4.4.1), during the initial analysis phase it became evident that what was initially perceived as a unique data set in fact contained three separate homogeneous datasets (homogeneous in the sense that the type of data was consistent in each set). Studies 2 and 3, on the other hand, generated one coding framework each.

*Table 5. Summary of the coding methods used across the three studies.*

	<b>Study 1 (chapter 5)</b>			<b>Study 2 (chapter 6)</b>	<b>Study 3 (chapter 7)</b>
<b>Coding Frame #</b>	<b>1.1</b>	<b>1.2</b>	<b>1.3</b>	<b>2</b>	<b>3</b>
<b>Analytical goal</b>	Code NTSB safety recommendation letters based on concerns regarding the use of the MSAW as identified by the NTSB.	Code NTSB safety recommendation s, targeting the “controllers’ lack of response” concern, based on the changes they demanded to the FAA.	Code a subset of the NTSB and FAA letters, according to their motivation for mandating/rejecting two of the NTSB-requested changes.	Find evidences in a heterogeneous data set in order to (re)construct a specific organisational experience.	Distil common themes from a homogeneous data set, and compare how they stand in comparison to the emerging framework.
<b>Coding method</b>	<b>Taxonomic</b> (see p. 65)	<b>Taxonomic</b>	<b>Comparative</b> (see p.67)	<b>Historical-evidentiary</b> (see p.69)	<b>Subsuming</b> (see p. 74)
<b>Codes stand for</b>	Types of concerns about the MSAW identified by the NTSB.	Types of changes requested by the NTSB to address one of the previously identified types of concern (the “controllers’ lack of response” concern).	NTSB’s and FAA’s respective motivations for mandating/rejecting two of the NTSB requested changes.	Organisational events, decisions, and conditions relevant to understanding MSAW set up and the management of MSAW-related nuisance alerts.	Attitudes, critiques, and professional experiences of experts useful for refining and expanding the OPHAI framework.
<b>Outcome</b>	Typological/descriptive framework	Typological/descriptive framework	Comparative framework	Conceptual framework	Conceptual framework
<b>Data display</b>	Graph	Matrix	Matrix	Matrix	Matrix

### 4.5.1. Study 1

As the first study of an exploratory research project, Study 1 turned out to be relatively challenging. Therefore, an introduction is needed before the description of the coding procedures is presented. The first analytical stage of Study 1 consisted of downloading and printing the initial data corpus—i.e., the set of NTSB’s safety recommendations, recommendation letters, and the FAA’s response letters (as described in § 4.4.1.1). This data corpus was read and re-read, while taking memos, asking clarification questions about the data, cross-checking the questions with colleagues and research supervisors, and supplemental documentation (mostly accident reports) when needed. This familiarisation phase was rather challenging: the originality of the data corpus resulted in a lack of immediate clarity about how to analyse it. Eventually, three distinct (but related) data sets were identified, namely:

- *Dataset 1:* This data set included the complete set of the MSAW-related NTSB safety recommendation letters and recommendations issued between 1977 and 2006. The coding of this data set was necessary to distil the categories of problems or concerns with the operation of the MSAW system identified by the NTSB. One of the identified problems, the “controllers’ lack of response”, emerged as critical for the purpose of the study and was subjected to further analysis;
- *Dataset 2:* This data set contained two sets of data. The first subset included the retrieved NTSB safety recommendations targeting the “controllers’ lack of response” concern. The second subset included the correspondence exchanged between the two agencies that addressed the implementation of the recommendations targeting the “controllers’ lack of response” concern;
- *Dataset 3:* This data set included the retrieved correspondence exchanges between the NTSB and the FAA addressing two requested changes that were refused by the FAA.

The identification of three distinct datasets meant that the first-level coding of Study 1 included three cycles—each cycle focusing on one data set. (Consequently, Study 1 first-level coding can in fact be viewed as three sequential sub-studies). The analysis of these data sets produced a total of three coding frameworks, which formed the basis of Study 1 second-level coding.

## 4.5.1.1. First-level coding for Study 1

## Coding methods for coding frameworks 1.1 and 1.2

The first two data sets were coded by means of the same coding method: *taxonomic coding* (Spradley, 1980). Taxonomic coding originated in ethnographic research and aims to construct taxonomies of the concepts that are typical of a specific culture, or cultural domain. According to Spradley (1980), every culture has its own cultural domains, which consist of categorisation systems. One problem is that such domains are part of (local) tacit knowledge and hence are not immediately intelligible to external observers. It is the researcher's task to extract these classification systems in order to make them visible to outsiders. In this study, taxonomic coding was used to identify (i) what the NTSB diagnosed as types of MSAW issues or concerns (coding framework 1.1); and (ii) what remedies the agency suggested to solve these issues (coding framework 1.2).

*Coding framework 1.1*

The set of MSAW-related NTSB safety recommendation letters that were retrieved (dataset 1) was coded for MSAW-related flaws. In particular, each letter (i) described one or more safety-critical events (civil aviation incidents and/or accidents); (ii) diagnosed the operational flaw(s) that lead to that event(s); and (iii) then provided a description of some desirable corrective change(s). An example of a safety recommendation letter (and its safety recommendations) is reported under Appendix C. In this phase, coding focused on understanding what flaws the NTSB identified for the MSAW—what aspects of the alarm chain did not work as intended and hence contributed, according to the NTSB, to the occurrence of a particular (set of) accident(s) and/or incident(s). When identified, relevant passages were assigned an initial In Vivo code. An example of a text passage coded as a type of NTSB diagnosed problem with the MSAW—lack of alarm activation—is reported in the table below.

**Table 6.** Coding framework 1.1: example of a code.

<b>Original text</b> (Portion of an NTSB safety recommendation letter)	<b>Code label</b> (Concern identified by the NTSB)
<p><i>“Because USAir 105’s premature descent...was inside the inhibit area designed to minimise the number of false alerts of low [altitude] aircraft that are maintaining a proper descent, the MSAW did not activate. The Safety Board believes that MSAW parameters can be adjusted provide for increased protection in areas where MSAW warnings are currently inhibited”</i> (NTSB safety Recommendation Letter 28 sept 1990, p. 4).</p>	<ul style="list-style-type: none"> <li>• Lack of MSAW activation</li> </ul>

After all the material had been read, the codes were listed onto a separate MS Word file. Once the list had been assembled, the coded text segments were systematically compared, checked for similarities and differences, sorted and re-arranged—the ultimate purpose being that of clustering them into higher-order categories that could stand for these segments. This process involved several iterations with the data. It was repeated until a stable definition of the categories of the framework was defined. The resulting framework consists of a relatively flat classification scheme made of descriptive categories. Each category crystallised types of identified NTSB concerns. The framework is reported in section 5.3.1.

### *Coding framework 1.2*

The same taxonomic coding approach was also applied to the second data set: the set of MSAW-related safety recommendations addressing the “controllers’ lack of response” concern. These recommendations were coded according to the nature of the specific corrective actions that the NTSB requested the FAA to implement in order to address this concern.

This analytical stage was necessary to build more precision and integrity into the analysis—that is, to distil the specific corrective changes embedded in the safety recommendations addressing the concern in question. In fact, some of these recommendations, although issued at subsequent points in time, essentially recommended the same corrective action. In other words, they repeated a request for a change that was already made in an earlier recommendation. Other recommendations, on the other hand, incorporated more than one corrective action. In other words the text of the recommendation contained requests for multiple changes. For instance, safety recommendation A-81-134 in Table 7 contains two design changes: a request to distinguish between the aural alarms of the MSAW and those of another alarm, the conflict alert; and a request to direct the MSAW aural alert only to the controller concerned.



**Table 7.** Framework 1.2: example of two codes.

<b>Quotation</b> (NTSB Safety Recommendation)	<b>Code labels</b> (Requested changes)
<p><i>“Redesign the low altitude/conflict alert at ARTS III facilities so that the audio signal associated with the low altitude alert <u>is readily distinguishable from that associated with the conflict alert</u> and <u>heard only by controller immediately concerned with the involved aircraft.</u>” (A-81-134)</i></p>	<ul style="list-style-type: none"> <li>• Distinguish between the MSAW and CA aural alert</li> <li>• Direct alert only to the concerned controller</li> </ul>

The safety recommendation analysis replicated the steps already described for coding framework 1.1, and resulted in a two-level descriptive coding framework. This includes eight types of *requested changes* (RC) (the lower-level categories), which were grouped under three types of *areas of change* (AOC) (the higher-level categories). Note that compared to coding framework 1.1, another analytical step was included in the development of coding framework 1.2. The data was assembled on a matrix in order to organise the identified requested changes over time. In particular, the matrix distributed the NTSB-requested changes according to the safety recommendation letters in which these changes were recommended. This analysis made it possible to trace the requests issued by the NTSB across the period studied. (The resulting matrix is reported under section 5.3.2.)

#### Coding method for coding frameworks 1.3 and 1.4

The first-level coding of Study 1 was concluded with an analysis of two of the requested changes identified in the previous coding phase, regarding which the two agencies held contrasting perspectives. On the one hand the NTSB requested and justified these changes; on the other, the FAA opposed them and presented a non-implementation rationale. Thus the development of coding frameworks 1.3 and 1.4 served the purpose of eliciting the contrasting rationales presented by the two agencies regarding the changes in question.

The data informing this analysis included (i) the NTSB safety recommendation letters in which the changes in question were presented; (ii) the initial FAA response; and (iii) the ensuing correspondence, in which the implementation of the changes was debated by the NTSB and the FAA. (The list of the specific data informing this phase of the analysis appears in Tables 11 (§ 5.3.3.1), and 13 (§ 5.3.3.2).

This analysis was based on the *comparative coding* method (Schreier, 2012). This approach allowed the construction of a comparative coding frame—i.e., a frame that

compares and illustrates the differences contained in two data sources. These can be distinguished based on any dividing criterion relevant to the study, such as different time periods, different stakeholder groups, and, as in this case, different official organisational views. Comparative coding was implemented in three steps:

1. *Identifying relevant passages in the first sources.* The relevant NTSB safety recommendation letter and the follow-up letters were read line-by-line in order to identify passages in which the NTSB justified the change in question.
2. *Identifying relevant data segments in the second sources.* The FAA response letters to the first NTSB safety recommendation letter and to the subsequent NTSB-follow up letters were read in order to identify the parts of the text in which the FAA reported its non-implementation rationale.
3. *Identifying differences between the two sources.* This step examined the categories identified in the previous two phases with the aim of understanding the differences between the NTSB and the FAA letters. The goal was to develop a coding framework that allows these differences to be represented as clearly as possible. The result consists of two summary matrices that detail the implementation and non-implementation rationales of the agencies regarding the two changes in question. These matrices are presented and explained in Table 12 (§ 5.3.3.1) and Table 14 (§ 5.3.3.2).

### 4.5.1.2. Second-level coding for Study 1

The final part of the analysis consisted of linking the results of the three sub-studies to the original research objective, namely the development of a qualitative framework of organisational precursors to HAI issues. This step involved a comparison and integration of the insights generated by the three frameworks of Study 1. This process led to the identification of the first category of organisational precursor of the OPHAI. The outcome of this phase is summarized in Table 15 (§ 5.4.2).

## 4.5.2. Study 2

### 4.5.2.1. First-level coding for Study 2

Study 2's data analysis had a strong investigative focus: the analytical challenge was that of examining a heterogeneous body of data in order to reconstruct the most salient traits of a whole organisational experience, namely the experience of introducing and managing the MSAW and the associated nuisance alert problem.

Because of this orientation, this study did not mirror the basic coding approach commonly found in sociological studies: the objective of the analysis was not to search the data for similar types of events, or incidents, across a homogeneous data set, consisting of, for example, interviews or, as already discussed in the context of Study 1, documentary data. Rather, data analysis in Study 2 was analogous to the investigative work of historians, forensic experts, and accident investigators. Notably, these investigators share a common analytical challenge: that of analysing multiple present-day evidentiary sources in order to reconstruct or infer some events of the past. For this reason the coding approach used in this study can be characterised as *historical-evidentiary*.

Data were read in a line-by-line fashion in order to identify passages which captured aspects relevant to the purpose of the research: organisational facts, conditions and decisions that were relevant to understand how the organisations under study (mis) managed the MSAW-related nuisance alert problem (see Table 8). Whenever such passages were found, they were marked using the same approach described in Study 1 (§ 4.5.1.1).

**Table 8.** Coding framework 2: example of the lower-level codes.

Data	Initial codes
<p><i>"At the beginning of **** our system was dying, it was a rudimentary monochromatic monitor. We needed a more stable and reliable system, as we were experiencing breakdown on a daily basis. So the main focus was on the implementation of the new system to get a higher-level of Flight data Processing capability. We knew pretty well the level of automation that was needed. On the other hand safety nets were a side dish, we had very little knowledge on them, and they were just a little feature that was offered as part of a larger system"</i></p>	<ul style="list-style-type: none"> <li>• Legacy system in need of replacement</li> <li>• System breakdown experienced on daily basis</li> <li>• Main focus during adoption was on the radar processing system</li> <li>• Limited knowledge of safety nets</li> </ul>

After all the data was read, the codes were listed in two separate MS Word files, one for Alphasky and one for Deltasky. This is the stage in which code comparison, checking for

similarities and differences, and subsuming began. This process allowed to filter out the most relevant codes and to distil the most relevant *case-specific categories*—i.e., categories of events, conditions, and processes relevant to the understanding of how the nuisance alert problem was handled within the two organisations. The resulting coding framework is reported in the table on the next page, and is presented in section 6.3.

### 4.5.2.2. Second-level coding for Study 2

The final stage of the analysis consisted of a within-case comparison: a comparison of the case-specific categories that emerged from the previous coding phase, with the aim of checking for similarities, differences, and potentially relevant dimensions across the two cases. This phase allowed the case-specific categories to be subsumed under three more abstract and theoretically-oriented categories, i.e., categories expressed at a level of abstraction that makes them potentially relevant also to organisational contexts other than the ones under study. This phase of the analysis led to the development of the matrix reported under section 0The final analytical step consisted of comparing the emerging categories of the OPHAI framework (identified by Study 2) with the category identified in Study 1.

Table 9. Coding framework 2: example of case-categories (and related lower-level codes) for Alphasky and Deltasky.

Case-Categories	Low-level codes	Quotations
<b><i>ALPHASKY (Success case)</i></b>		
<b>A1. Having an operational need for introducing the MSAW</b>	<ul style="list-style-type: none"> <li>Implementing the alarm in response to a CFIT</li> <li>MSAW implementation perceived as high priority issue</li> </ul>	<p>“...a particular effort [has] to be made as soon as possible [by Alphasky] to complete the design and implementation by air traffic control services of a ground based system for detection of an aircraft in dangerous proximity to terrain, whenever technically possible” (Safety recommendation)</p> <p>“In ***[year omitted] there was a CFIT in *** [crash site omitted]. At that point the incorporation of a ground proximity alert function into the approaches became a high priority action.” (Head of ATC, site 1)</p>
<b>A2. Clarifying the operational role of the MSAW</b>	<ul style="list-style-type: none"> <li>First studies into the MSAW began prior to adoption</li> <li>Considering the MSAW role of “attention grabber”</li> <li>Choosing to develop the MSAW as “hazard resolver”</li> </ul>	<p>“It was necessary to wait until the mid-eighties for the first studies by [Alphasky] into the development of a real-time ground based proximity detection system. In 1988, the [internal R&amp;D centre] made the first survey of requirements expressed by the operational personnel and the [office of accident investigation]...Opinions differ from county to another regarding the nature of the MSAW function and the service it should provide. Even if everyone agrees about what an MSAW is and should only be a control aid, some feel that the alerts it generates should simply attract the controller’s attention, it then being up to him to analyse the situation and make the necessary decision...This viewpoint is not shared by ***, which has chosen to develop a reliable system in which all alerts are justified, and which should entail no situational analysis work by the controller, but rather a reflex action informing the aircraft concerned.” (Alphasky documentary source)</p>
<b>A3. Recognizing the problem of nuisance alerts</b>	<ul style="list-style-type: none"> <li>Acknowledging the detrimental problem of overexposure to frequent nuisance alerts</li> </ul>	<p>“The appearance of a large number of false alerts and undesirable alerts can be a factor in a certain dilution of controller vigilance and reduced confidence in the system, which in the end means that the controller-system structure no longer correctly provides the collision avoidance alert service.” (Alphasky documentary source)</p>
<b>A4. Parameterisation process</b>	<ul style="list-style-type: none"> <li>Recording a sample of live air traffic data</li> <li>Installation on local sites</li> <li>Feeding the alarm with real data</li> <li>Operational on-site testing</li> <li>Final approval</li> </ul>	<p>“I think [parameterisation] comes from recounting of real traffic and analyzing real traffic is and counting the number of alerts and deciding which one is useful which one is not. So they’ve taken a sample traffic studying the number of alerts, deciding this is useful this one is very useful, this one is nuisance. This is the way the system is parameterized.” (Alphasky safety net expert)</p> <p>“[when tuning the alarm] we collect traffic with a period of 1 month, 2 months.... Having real traffic data is key. Sometimes there are some days in which I do not have the recordings for doing the tuning. For ***[airport name omitted] I collect traffic for just one month, because the higher the traffic, the larger the sample of recorded traffic” (MSAW engineer)</p>
<b>A5. Positive attitude towards air traffic controllers’ involvement</b>	<ul style="list-style-type: none"> <li>Having experts controllers permanently allocated to MSAW parameterisation</li> <li>Engineers-Controller consultation</li> </ul>	<p>“We have at least two controllers in ***[site name omitted] who get involved in the set up the MSAW. They are expert controllers yes, they have operational experience, but they also are knowledgeable about how to work with the safety net engineers. Usually what they do is...well, they work with a plot like this one, for instance this one shows all the alerts generated by the system at [***airport name omitted] over a day [...] [the expert controller] tell the engineer which one is good and which is bad” (Alphasky safety net expert)</p>

<p><b>A6. Supporting software tools and organisational roles</b></p>	<ul style="list-style-type: none"> <li>• Supporting software tools</li> <li>• Supporting organizational roles</li> </ul>	<p><i>"To evaluate the real environment of the MSAW algorithms, the [Alphasky R&amp;D centre] produced a demonstrator to visualise the alerts on a radar picture and tools for recording, flight path analysis and statistics. The demonstrator was installed within [the R&amp;D facilities] (in 1993) and in ***[sites name omitted] (1994 and early 1995). The purpose of the evaluations was to improve the MSAW algorithms and to check their reliability, their efficiency, and to optimise the operating parameters". (Alphasky documentary source)</i></p> <p><i>"Currently the parameterisation and improvement of the system involves two safety net experts, the validation team, and a committee that reunites once or twice a month...plus there is national leader, since it has been nominated things have improved a lot, as the issues with the system get addressed sooner." (Alphasky safety net expert)</i></p>
<p><b>A7. Specifying the requirements to the manufacturer</b></p>	<ul style="list-style-type: none"> <li>• Having established a long lasting, R&amp;D, relationships with the manufacturer</li> <li>• Specifying requirements to the manufacturer</li> <li>• Software modifications done by the design and validation team</li> <li>• The design and validation team acting as a point of contact with the manufacturer</li> </ul>	<p><i>"Concerning MSAW...[Alphasky] did most of the work. They did the concept development, the high-level and the low level design. They have a very competent organisation which can cover the entire lifecycle. In this particular case, they took the MSAW of ***[manufacturer name omitted] apart, modified it, and improved for a few years, and then give it back to industry a few years later. Today the MSAW sold by ***[manufacturer name omitted] is a best in class." (EUROCONTROL Spin leader)</i></p> <p><i>"[the specification of MSAW requirements] was to be followed by the development – using a software developed by ***[manufacturer name omitted]...of a software incorporating the MSAW and the AIW functions, and improving the initial software by two major changes which are on the one hand the incorporation of [design detail omitted] with a very fine mesh, thereby improving alert precision and quality, and on the other the detailed parameter configuration of the airport areas" (Alphasky documentary source)</i></p> <p><i>"All software modifications are done in ***[site name omitted] by the Design and Validation team. When something with the MSAW does not works, a report is sent to this team, which processes all reports from [Alphasky] sites, and then generates the requirements. These are then passed to [the manufacturer] for implementation. Once this is process is done, a CD with the update is sent to us, where the technician install the software on the local radar. No modification is made by the local facility, and no code is written/modified locally – this is actually forbidden for us. This stuff is done by the Design and Validation Team, which also is the single point of contact with the manufacturer"(Alphasky safety nets expert)</i></p>
<p><b><u>DELTASKY (Less successful case)</u></b></p>		
<p><b>D1. Lacking clarity about the operational need for introducing the MSAW</b></p>	<ul style="list-style-type: none"> <li>• The original radar processing system was dying</li> <li>• Initiated modernization effort in 2003</li> </ul>	<p><i>"By 2003 ***'s previous system (deployed in early 80's) was way over its lifecycle and an immediate replacement was needed. In 2003 [Deltasky] entered a contract with ***[manufacturer name omitted] for a new Data Processing System (Integrated FDPS and RDPS) to serve as an interim solution, while the work of specifying operational requirements for an overall modernization programme was kicked off simultaneously. The work with ***[manufacturer name omitted] was carried out under the ***[modernization programme name omitted] [...which provides] for STCA, MSAW and APW (or RAI – Restricted Area Intrusion Warning, as referred to by [the manufacturer], and has been operational in ACC ***[name of te omitted] since 2004, and still is."(Deltasky project manager)</i></p> <p><i>"Our burning issues was to have a stable [radar processing and flight data] system. The old one was</i></p>

	<ul style="list-style-type: none"> <li>• Main priority was to have a stable system in place</li> <li>• Knowing the required functionalities of the new system</li> <li>• Purchasing the MSAW as a COTS</li> </ul>	<p><i>unstable, dying in our hands virtually...So our burning issue was to have stable system based on new pc platform and flat displays, good flight plan processing...We did not have any external driver for implementing MSAW and other safety nets; we needed to catch up with industry standards.” (Deltasky project manager)</i></p> <p><i>“MSAW came in a package of safety nets in ***[manufacturer name and the supplied system omitted]...Brought into operations on *** May 2004.” (Deltasky power point presentation)</i></p>
<b>D2. Lacking clarity about the operational role of the MSAW</b>	<ul style="list-style-type: none"> <li>• MSAW seen as a minor technical system</li> <li>• Lacking awareness about MSAW</li> <li>• Learning the operational purpose of MSAW after implementation</li> </ul>	<p><i>“On the other hand safety nets were a side dish, we had very little knowledge on them, and they were just a little feature that was offered by ***[manufacturer’s name omitted] as a part of a larger system. Also we were not aware of the necessity of having safety nets. This is reflected in the way we parameterize safety nets later on” (Deltasky project manager)</i></p> <p><i>“Lesson learned from SPIN [after implementation]: Prevention of CFIT [is] the sole purpose of MSAW” (Deltasky power point presentation)</i></p>
<b>D3. Purchasing the system in the absence of specified requirements</b>	<ul style="list-style-type: none"> <li>• Lack of MSAW requirements</li> <li>• Trusting the manufacturer</li> <li>• Relying on the manufacturer for the system parameterisation</li> </ul>	<p><i>“Well, you may consider that only the MSAW acronym appeared in the contract, with no dedicated requirement for this system. We accepted to implement it after they told us they could implement also this system [...] We trusted the manufacturer very much, we selected it among one of the best in the industry, we knew they had implemented the same systems on another ANSP (I know because I visited them), so we assumed we could rely on them also for the setup of the MSAW and safety nets...” (Deltasky project manager)</i></p>
<b>D4. Manufacturer providing a sub-optimal alarm</b>	<ul style="list-style-type: none"> <li>• Manufacturer considering the supplied grid adequate to the local terrain</li> <li>• Grid too unrefined for the specific terrain</li> </ul>	<p><i>“The engineers of ***[manufacturer name omitted] deemed the [specified] grid to be ok. However, when we started using it, we realised it was so unrefined. If you cannot refine your warning altitude to distinguish between an aircraft that is on a correct course of approach from one that is not, then you have a nuisance.” (Deltasky IT engineer)</i></p> <p><i>The grid supplied by the manufacturer consisted of a 10 by 10 square made by modules of 8x8NM, something that is far too broad for the MSAW set up (Deltasky power point presentation)</i></p>
<b>D5. Realising the nuisance alert problem after implementation</b>	<ul style="list-style-type: none"> <li>• Alarm generating too many nuisance alerts</li> <li>• Alarm lacking a single click inhibition function</li> </ul>	<p><i>“The alarm generated an alert for every single approaching aircraft” (Controller)</i></p> <p><i>“Feedback from operations [was] Turn it off!” (Deltasky power point presentation)</i></p> <p><i>“The alarm did not contain a single-click inhibition function. Including a single click for alert inhibition is part of the lessons learnt following the first MSAW implementation.” (Deltasky engineer)</i></p>
<b>D6. Realising the need for importing parameterisation expertise</b>	<ul style="list-style-type: none"> <li>• Lacking the skill set to parametrize the system</li> <li>• Seeking “help” outside the organisation</li> </ul>	<p><i>“It was when the manufacturer had packed its suitcase and gone that we realised we did not have access to system functions, the knowledge, the tools, and the competence to do the remaining implementation process on our own [...] This is also the stage we realised we needed external help, so we turned to SPIN, this is where we really found a way to get to know these system, their purpose, and how they have to be set up.” (Deltasky project manager)</i></p>

### 4.5.3. Study 3

Study 3's data analysis consisted of extracting, from a homogeneous data set made of expert account, facts and attitudes relevant for corroborating the categories of the emerging OPHAI framework, defined in Study 2.

*Subsumption* (Mayring, 2010) was the coding approach adopted in this study. It consists of coding data with reference to an existing coding framework or classification system. This was in fact the first study of the present research in which data analysis was driven by an existing theoretical framework, namely the initial version of the emerging OPHAI framework. Coding proceeded by means of close reading of the text. Whenever a passage that was deemed to fall into one of the existing categories of the framework was found, this was marked by assigning it a code according to the theoretical category of the framework, i.e., OP1, OP2, and OP3. In this way, the codes assigned were instrumental to the conceptual subsumption of parts of the data to the three categories of identified precursors. The relevant coded text passages consisted of the SMEs' past professional experiences, and personal attitudes towards the categories of the framework.

Note that although code development was driven by an existing classification system, i.e., the initial version of the framework, parts of the text that did not fit into the original framework were not ignored. These passages were marked as well, as they also could contain useful ideas for the expansion or refinement of the initial framework. Most importantly, having code development driven by an existing classification system did not mean that coding proceeded mechanically in a strict top-down fashion. This is the case because the categories of the initial framework were conceptually-oriented: they allowed coded textual data to be sorted only into three high-level categories, which however did not account for lower-level differences present in the data in each category. Therefore, within each category, a lower-level classification scheme had to be developed in order to account for these differences. And developing such scheme required the researcher to work inductively from the bottom—beginning with the data—while considering the pre-existing categories.

After coding the data set, the coded segments were listed in a separate MS Word file. At this point, these segments were examined in terms of (i) their differences and similarities between the incidents and events they contained; and (ii) their relevance to the initial OPHAI's categories—i.e., the extent to which the coded category supported/refuted each specific category. The resulting coding framework is presented in section 7.2.



#### 4.6. WITHIN-STUDY VALIDATION STRATEGIES

---

In the context of this research it was not feasible to have the data coded by two different people because of resource constraints. Thus, in all the three studies, the coding process was done by the researcher alone. Therefore, to guard against the researcher's own bias, the following strategies were used within each study (in addition to the validation strategies already described in § 4.3.4).

1. **Review of the coding framework at different points in time.** The respective coding frameworks developed in the three studies were each reviewed for consistency at least three times, with no less than a month elapsing between reviews. During each review, the author qualitatively verified the ability of the framework's categories to adequately represent the data for which the categories stood.
2. **Peer debriefing.** This strategy consisted of checking the accuracy of intermediate and final results with individuals familiar with the research. This strategy was implemented in different ways for Studies 1 and 2. In Studies 1 and 2, the frameworks' categories were constantly verified with a supervisor knowledgeable about the ATM domain and automation. In addition, intermediate results of both studies were presented to both research supervisors and members of the EUROCONTROL Safety Team that were familiar with the research. The results of these studies were also disseminated at two academic conferences (Studies 1 and 2) and in a journal article (Study 1). (See Appendix F for the list of work published in this research.) Also Study 3 results were verified by one supervisor knowledgeable about automation and ATM.
3. **Member checking.** In Study 2, visit reports following site visits were sent back to the research participants within three days for cross-checking purposes. In Study 3, collection of the qualitative questionnaires was usually followed up with further questions sent via e-mail to research participants and aimed at verifying/expanding on some of the issues reported in the questionnaire.
4. **Data triangulation.** Study 2 made use of different types of data sources. Thus, an important part of the analysis consisted of triangulating or cross-checking these sources.
5. **Prolonged engagement in the domain.** This strategy was especially relevant for Study 3. The attendance of three SPIN meetings during Study 2 facilitated access to Alphasky and Deltasky, and also ensured that by the time Study 3 was executed the researcher was known to members of the SPIN community. Notably, engagement with the

community of research participants usually creates trust, and establishes rapport thus making it easier for participants to disclose relevant information.

#### **4.7. CHAPTER CONCLUSIONS**

---

This chapter presented the research strategy adopted in the present thesis. The chapter initially provided an overview of the three studies used to develop the OPHAI framework: two historical qualitative case studies centred on the organisational trajectory of a known HAI issue, the MSAW-related nuisance alert, followed by a third corroboratory study. The chapter explained the theoretical foundations of this strategy. It also determined that such a strategy is adequate to the present research, as it is consistent with its exploratory purpose, it is feasible, and it mitigates satisfactorily the major validation issues typical of single case study designs. Furthermore, the chapter described the methodological aspects related to data collection and data analysis procedures, and the qualitative validation strategies that were used in the three studies. Having completed the description of the research's strategic and methodological aspects, the following three chapters present the results of the three studies.

# Chapter 5.

## STUDY 1 Results

### 5.1. CHAPTER INTRODUCTION

---

This chapter reports on the first empirical study of this thesis: an explorative and retrospective case study that inquired into the implementation and improvement history of the MSAW in the US. In this context, the study aimed to explore what organisational conditions occurring at the blunt end of the system influenced the handling of the MSAW nuisance alarm problem. The case is based on the analysis of (i) safety recommendation letters and recommendations targeting the MSAW, as issued by the National Transportation Safety Board (NTSB) to the Federal Aviation Administration (FAA); the analysis of (ii) the FAA's initial response, and (iii) the ensuing correspondence letters in which the implementation of these recommendations was debated. The chapter is organized as follows:

- Section 5.2. provides background information useful for understanding the case study: a brief overview of the US MSAW (§ 5.2.1), and a description of the institutional cycle through which the documents analysed in the study are usually developed, generated and debated between the NTSB and the FAA (§ 5.2.2).
- Section 5.3. presents the results of Study 1 first-level coding, i.e., four coding frameworks;
- Section 5.4. presents the results of Study 1 second-level coding, i.e., a coding framework that builds on the first-level frameworks in order to identify which organisational precursor/s was/were at play in the present case. In so doing, this section provides an initial version of the framework of the organisational precursors to human automation interaction issues (OPHAI).

## 5.2. BACKGROUND

---

### 5.2.1. US MSAW

The specific call for introducing the MSAW system in the US occurred following the crash of Eastern Air Lines flight 401 while on approach to Miami International Airport in December 1972 (NTSB, 2000, 1973). Following this accident, the NTSB issued safety recommendation A-73-46, which demanded the FAA to equip their radars with a low altitude warning.

in 1974, after an initial consultation phase, the FAA contracted Sperry Rand's Univac Division to develop an add-on for the radar system then in use in terminal facilities, the ARTS III (FAA, No date). The first MSAW system was commissioned at Los Angeles airport in 1976, and by 1977 the alarm was operational at all 63 major US airports equipped with ARTS III (FAA, No date). In 1981, the FAA also introduced a version of the MSAW for en-route control, called E-MSAW. This new feature was operational at 14 of the 22 air route traffic control centres by 1987 (FAA, No date).

Despite its adoption by various FAA control centres, the MSAW has been the target of several safety recommendations issued by NTSB to the FAA that aimed at correcting some criticalities with the alarm. This study analyses (i) these MSAW-related safety recommendations; (ii) the safety letters conveying these recommendations; (iii) the corresponding FAA's response; and (iv) the ensuing correspondence exchange letters by means of which the NTSB either monitored the implementation of a given recommendation, or reiterated the need for implementing one.

### 5.2.2. The organisational context of analysis

The investigation of aviation accidents by independent bodies is one of the most sophisticated forms of present-day safety learning and improvement. The core of this process consists of identifying the probable causes of a safety event, be it an accident or an incident, and producing safety recommendations aimed at preventing its re-occurrence. In the US, the NTSB is the independent federal agency charged by the US Congress to investigate every civil aviation accident, as well as significant accidents occurred in other transportation modes—railroad, marine, and pipeline (NTSB, 2004a). Since its foundation in 1967, the NTSB has investigated more than 124,000 aviation accidents and issued more than 12,000 safety recommendations, both within and outside

aviation (NTSB, 2004a).

NTSB safety recommendations are the most important part of the NTSB mandate (NTSB, 2004c). They address the need for remedial actions that usually emerge during or upon the completion of an NTSB investigations. Safety recommendations can in fact be issued before the completion of an investigation, furthermore supporting evidence can also come from fact-finding reviews and safety studies. When a safety recommendation is defined, the NTSB defines its status as “open” and communicates it to the organisation best able to act on the problem, whether public or private (NTSB, 2004b), through a safety recommendation letter. Each letter can contain one or more safety recommendation(s) depending on the criticalities identified by the NTSB. These safety recommendations are issued without performing a cost-benefit analysis. An example of a safety recommendation letter and its safety recommendations can be found in Appendix C.

#### 5.2.2.1. FAA Response to NTSB Safety Recommendations

On receipt of an NTSB safety recommendation letter, the FAA assigns the recommendation(s) contained in the letter to its Office of Accident Investigation (AAI). The AAI is the FAA body responsible for the processing, monitoring and management safety recommendations. First, it reviews the safety recommendation to verify its adequacy, accuracy and appropriateness in resolving the safety issue addressed. Second, it assesses the feasibility of implementing the safety recommendation or an alternative action (DOT/FAA, 1995). This is necessary because, as noted above, NTSB safety recommendations are developed without performing a cost-benefit analysis. Finally, after evaluation, the AAI prepares the initial official response to be issued to the NTSB, usually within 90 days of the receipt of the safety recommendation letter (DOT/FAA, 1995).

#### 5.2.2.2. Ensuing correspondence exchange

After issuing the safety recommendation to the FAA, the NTSB defines the safety recommendation status as “open”. The FAA responds by presenting their position with regard to the recommended action, and their planned response, if any. The FAA may in fact decide not to implement any response action, regardless of how urgent the NTSB deems it to be, since the latter does not have any regulatory power, and can only reiterate its request(s) to adopt the recommendation (e.g., Danko, 2010; Carlisle, 2000).

This first correspondence exchange may be followed by the NTSB issuing a second letter

to persuade the FAA to implement a given recommendation, or to demand an update about its implementation status. Several written exchanges between the two organisations may be necessary between the two organisations to debate about acceptance and implementation status of a given safety recommendation. This process ends with the closure of the safety recommendation by the NTSB, which can close the recommendation by assigning it the status of “Acceptable action”, “Acceptable, alternative action”, or “Unacceptable, action”. This completes the description of the context of analysis. The next section presents Study 1’s results.

### 5.3. FIRST-LEVEL CODING RESULTS

---

This section reports on the results of the first-level coding phase of Study 1. As described earlier (§ 4.4.1), because three distinct (though related) data sets were used, Study 1 first-level coding consisted of three coding cycles. The three datasets are as follows:

1. The set of the retrieved NTSB safety recommendation letters. These were analysed in order to identify the problems, or concerns, that the NTSB had found with the MSAW since the alarm entered into service;
2. The set of the retrieved NTSB safety recommendations targeting a specific NTSB concern (the “controllers’ lack of response” concern). Resulting as the most problematic from the previous analysis, this concern was analysed to understand the remedial changes requested, by means of dedicated safety recommendations, by the NTSB;
3. The set of the retrieved NTSB safety recommendation letters and the associated correspondence exchange letters between the NTSB and the FAA concerning two requested changes, regarding which the two agencies developed two opposing rationales. This dataset permitted investigation of the sources of disagreement between the two agencies.

The analysis of these datasets produced a total of four coding frameworks, because the third dataset’s analysis involved two coding cycles: one for each change requested. These coding frameworks are described next.

#### 5.3.1. Coding framework 1.1: MSAW-related concerns identified by NTSB


Study 1’s first (preliminary) analytical step consisted of coding the NTSB safety recommendation letters retrieved depending on the MSAW-related concern(s) they raised: in each letter the concern corresponds to passages of text that answer the question “Based on the evidence—i.e., accident(s) and/or incident(s)—reported in this specific letter, what main problem(s) does the NTSB sees in relation to the MSAW?”. These concerns capture the NTSB’s view about the ways in which, in the relevant accident(s) and/or incidents(s), the joint work unit composed of the MSAW and the controller failed to function in the intended way—i.e., failed to deliver a low altitude warning alert to the concerned aircrew. Essentially, concerns correspond to conditions (identified by the NTSB) that jeopardised the alarm’s protective potential.



The analysis revealed that the NTSB has identified nine categories of MSAW-related concerns over the period 1977–2006. These were:

- **1. Controllers' lack of response.** This category groups all of those safety letters reporting on the controllers' tendency to ignore MSAW warnings, even in the presence of true, reliable alerts;
- **2. Inhibited MSAW processing.** This concern emerged in one US airport where the radar system deactivated MSAW processing whenever it received low quality radar return signals, which were presumably due to reflection signals generated by tall buildings near the runway;
- **3. MSAW System not installed.** This concern refers to the MSAW not being installed in remote sites where a particular accident/incident occurred. Because of this, the controllers of these sites, were unable to issue low altitude alerts;
- **4. Permanent inhibition for Visual Flight Rules (VFR) flights.** MSAW inhibition for aircraft associated to VFR computer codes was a normal condition during routine operations to reduce nuisance alerts: VFR flights routinely typically fly lower than Instrument Flight Rules (IFR) flights. However, this permanent MSAW inhibition emerged as critical when VFR flights were in an emergency situation. The risk was that VFR went unnoticed or received little help from ATC;
- **5. Over-dimensioned inhibited area.** Over-dimensioned inhibited areas resulted in the MSAW not detecting aircraft flying below the minimum safe altitude, because the descent would occur in areas where the alerting function was permanently inhibited to minimise unwanted alerts;
- **6. Misplaced capture box.** This category refers to the mistaken misplacement of the MSAW capture box, or grid, with regard to its intended position. Such a misplacement resulted on the MSAW not detecting aircraft flying below the minimum safe altitude.
- **7. MSAW speaker not installed.** The absence of the speaker in tower control centres was considered as critical by the NTSB, because in these settings controllers, in order to control traffic movement, spend most of their times looking out of the window at aircraft. Thus, controllers' attention can be redirected more effectively by means of an aural warning;
- **8. Lack of alarm integrity overseeing.** This concern emerged when an MSAW audio speaker was found covered with masking tape as a means to silence it. It refers to the (in)ability of the organisation to monitor the integrity of the system;

- 9. Inconsistent alarm sensitivity across sectors.** This category refers to the presence of inconsistent MSAW settings between approach and tower control centres, so that the same risk of Controlled Flight into Terrain (CFIT) accidents was not consistently alarmed across different sectors.



**Figure 7.** Categories of concerns regarding the MSAW as identified by the NTSB over the period of analysis.

The graph in Figure 7 relates the concerns identified to the number of NTSB safety recommendation letters in which the concerns were mentioned. The graph indicates that the concern n. 1, the “controllers’ lack of response” concern, emerged as a long-standing issue: it was in fact repeatedly mentioned in five NTSB safety recommendation letters. Each of the other categories of NTSB concerns was addressed in one letter only. This difference is due to the fact that the solution to these latter concerns was more straightforward, as they primarily reflected relatively “simple” technological problems: all

had to be done to solve the problem was installing a missing technological component or modifying an existing one. On the other hand, the solution to the “controllers’ lack of response” concern was not so straightforward: the problem in this case was not solely technological, but also had a human component—the air traffic controller’s behaviour triggered in response to the MSAW alert. Another important observation to make at this point is the fact that the nuisance alert problem did not surface as an independent category of concern. As it will be discussed next, this problem was debated in the context of the “controllers’ lack of response” concern.

### **5.3.2. Coding framework 1.2: Areas of changes requested by the NTSB to address the “controller lack-of-response” concern**

A subsequent analysis was conducted to identify which changes were asked by the NTSB to the FAA in order to address the “controller lack-of-response” concern. This analysis classified the NTSB safety recommendations that were issued in order to solve this concern depending on the nature of change they requested. The analysis identified three areas of change: (i) human machine interface design, (ii) procedure, and (iii) nuisance reduction.

1. *Human Machine Interface (HMI) design.* This area of change includes 5 requested changes (RC), corresponding to 10 safety recommendations, which demanded changes to be made to the alarm design to improve its visual and audio presentation to air traffic controllers. In particular, the NTSB issued the following requested changes:

- One requested change, RC1, aimed at improving the informativeness of the MSAW aural alert. This change was motivated by the fact that the MSAW aural alert was identical to that of another alarm, the conflict alert or CA—an alarm system intended to warn controllers of the risk of mid-air collisions<sup>6</sup> (MACs). Thus, the same aural tone could signify the risk of a CFIT or that of a MAC. For this reason the NTSB wanted to differentiate the aural alerts of the two alarms;
- One requested change, RC2, aimed at making the MSAW flashing rate more visually appealing compared to that of the data block—the block appearing on the radar display containing flight’s alphanumeric information—during transfer between different sectors;

---

<sup>6</sup> The CA is in fact the US equivalent of the Short Term Conflict Alert (STCA), which is used in Europe.

- Two requested changes, RC3 and RC4, aimed at decreasing the exposure to alerts generated by other working positions in the same control room. In particular, RC3 aimed to direct the alarm only to the controller with direct control of the aircraft that activated the MSAW; RC4 included the request to incorporate an alarm inhibition function, so to silence MSAW alerts activated by other working positions;
  - One final requested change, RC5, requesting the total redesign of the HMI to make the alert more attention capturing. In the fifth and final safety recommendation letter reporting on the lack-of-response concern, the NTSB questioned the ability of the entire MSAW alerting method to reliably capture the attention of air traffic controllers.
2. *Procedure*. This area of change includes two corrective changes—RC6 and RC7, corresponding to two safety recommendations—that demanded the FAA changes to existing procedures to improve the controller’s response to the MSAW. RC6 asked the FAA to include the shift supervisor in the alerting loop, so as to ensure that whenever the MSAW was activated s/he could intervene to ensure that that the relevant controller responded appropriately. RC7 demanded regulatory changes to make it mandatory for air traffic controllers to transmit the MSAW alerts to the relevant pilot whenever the aural alert was heard.
  3. *Nuisance alert reduction*. This area of change consists of one requested change, RC8—corresponding to one safety recommendation that asked the FAA to reduce the rate of unwarranted nuisance alerts in control centres, because the NTSB had noted that controllers were exposed to such alerts on a routine basis.

The matrix reported in Table 10 provides a chronological overview of (i) the areas of change and (ii) the changes requested by the NTSB to address the lack-of-response problem. The matrix distributes the NTSB requested changes, grouped by area of change, across the five safety recommendation letters in which they were issued. The cells in the matrix show the identifier of the safety recommendation that carried a given change. Identifiers in bold signal safety recommendations that have triggered a non-implementation rationale on the part of the FAA.

The matrix shows that the first area of change, “HMI design”, is both (i) the older and (ii) the denser area of change. Safety recommendations belonging to this area were issued first in 1981 (n=2), then in 1984 (n=1), in 1990 (n=3), in 1997 (3), and lastly in 2006 (n=1).

This makes a total of 10 safety recommendations (=77% of the total) issued by the NTSB in relation to this area of change. The “HMI design” area is followed by the second improvement area, “Procedure”, which includes two corrective actions issued in 1984 and 1997 by means of two safety recommendations (=15%). The third area, “nuisance alert reduction”, is instead the sparser and younger area, including only one requested change issued by means of one safety recommendation in 2006. As an intermediate conclusion, it can be said that, over the period 1981–2006, the NTSB viewed the “controller lack-of-response” concern mostly as an issue of improving the HMI Design and the procedural response to the MSAW. It is in 2006, that the NTSB considered the importance of tackling the problem of nuisance alerts as a stand-alone problem.

**Table 10.** Matrix of the changes requested by the NTSB to the FAA.

NTSB Safety Recommendation Letters		1st, Oct 6, 1981	2nd, Aug 13, 1984	3rd, Sept 28, 1990	4th, Apr 16, 1997	5th, July 12, 2006
Areas of Change	Requested Change (RC)					
<b>1. HMI Design</b>	1. Disambiguate between MSAW and CA aural alerts	<b>A-81-134*</b>	<b>A-84-83</b>	<b>A-90-161</b>	-	-
	2. Make MSAW flashing at a higher rate than the data block during transfer	A-81-135	-	A-90-163	A-97-25	-
	3. Direct MSAW aural alarm to the relevant controller only	<i>Included under A-81-134</i>	<i>Included under A-84-83</i>	A-90-162	A-97-26	-
	4. Include possibility to inhibit the alarm	-	-	-	A-97-26	-
	5. Completely redesign MSAW alerting method	-	-	-	-	A-06-44
<b>2. Procedure</b>	6. Include supervisor in the alerting loop	-	-	-	A-97-27	-
	7. Amend regulations to make mandatory the transmission of MSAW alert to pilots	-	<b>A-84-84</b>	-	-	-
<b>3. Nuisance Alert reduction</b>	8. Reduce exposure to nuisance alerts	-	-	-	-	A-06-45
<b>LEGEND</b> Identifiers in the cells correspond to safety recommendations; *= Safety recommendations in bold have been dismissed by the FAA.						

Also, the matrix suggests that the first and second areas of change, i.e., “HMI design” and “Procedure”, include some changes that were resisted by the FAA. In particular, the FAA expressed a non-implementation rationale for RC1, “Disambiguate between the MSAW and the CA aural alert”, and RC7, “Amend regulations to mandate the transmission of MSAW alert to the pilot”. The different positions of the NTSB and the FAA regarding these changes are qualitatively analysed in the next section. They contain important insights for the understanding of the history of the nuisance alert problem.

### 5.3.3. Coding Framework 1.3 and 1.4: analysis of the contrasted changes

This section reports on the NTSB’s and FAA’s positions regarding the two contested changes, mandated by the NTSB to solve the “controllers’ lack of response” concern, that were not accepted by the FAA, namely RC1 and RC7.

#### 5.3.3.1. RC1: Disambiguate between the MSAW and CA aural alerts

The NTSB issued three times a request to have two different aural alert tones to be used for the MSAW and the CA. Since 1981, the NTSB noted that the same aural alert was used for signalling two different risks, CFIT and MAC. This change was issued in 1981 as safety recommendation A-81-134; in 1984 as A-84-83; and then in 1990 as A-90-161 (see Table 11). Following each of these requests, a written exchange between the two agencies ensued, in which the FAA stated (and restated) its non-implementation rationale for the change in question, and the NTSB restated the importance of implementing it. The respective rationales of the two agencies identified in the analysis are summarised in Table 12 and explained in the following two sections.

**Table 11.** Safety recommendations containing the request to differentiate between the MSAW and CA aural alerts: for each recommendation, the table reports the NTSB safety letter in which the recommendation was made, and the ensuing letters exchanged by the NTSB and FAA.

Safety recommendation	Issued in NTSB safety recommendation letter:	Ensnuing letters exchanged between the FAA and NTSB
A-81-134	A-81-132 through -138, Oct 6, 1981	(1) FAA Letter Dec 21, 1981 (2) NTSB letter July 8, 1982 (3) FAA Letter Sept 1, 1982 (4) NTSB Letter March 8, 83
A-84-83	A-84-82 through -84, Aug 13, 84	(1) FAA Letter Jan 18, 1985 (2) NTSB July 1, 1985 (3) FAA Letter Aug 26, 85 (4) NTSB Letter Jan 24, 1986
A-90-161	A-90-161 through -163,	(1) FAA letter Jan 18, 91

	Oct 29, 1990	(2)	NTSB letter May 20, 1991
		(3)	FAA letter Dec 28, 1992
		(4)	NTSB letter April 16, 93

**Table 12.** Coding framework 1.3: summary of the NTSB’s and FAA’s positions over NTSB Requested Change 1.

<b>NTSB Rationale</b>	<b>FAA Non-implementation rationale</b>
1-Controllers are continually exposed to CA and MSAW nuisance alerts; 2-This continued exposure conditions them to frequently dismiss the alarm; 3-However, this may result in the controller ignoring a reliable MSAW alarm; 4-A perceptual reinforcement of the alarm signal is needed to better attract the controller’s attention.	1-There is no need to differentiate between the CA and MSAW aural alarms; 2-The aural alarm serves as an attention getter or general warning: it means “scan the display”; 3-The controller does not have to take a control action based on the aural alert alone; 4-The visual alert serves as a specific warning.

**NTSB’s position: reinforcing the alarm’s ability to attract controllers’ attention**

The NTSB maintained that controllers’ response to the MSAW could be improved by differentiating the MSAW and CA aural tones: the two alarms used to have the same aural tone, in fact, and were presented in an operational environment filled with nuisance or unnecessary alarms, which conditioned controllers to ignore aural warnings; as a result, modifying the HMI design in order to disambiguate between MSAW and CA aural tones—i.e., making them more representative of the specific hazard they denote—was seen as a necessary improvement to increase the alarms’ potential to attract controllers’ attention.

The NTSB’s acknowledgment of the over-exposure to nuisance alarms, the consequential controller’s habituation and desensitisation, and the need to change the MSAW design so as to make the aural alarm more informative can be found in several passages:

*“This situation and the others mentioned above [common tone and sources for LA and CA, and alarms which have their origin in another controller’s airspace] result in repetitive alerts which, in turn, condition the controller to dismiss the alarms or alerts! (i.e., “the cry” wolf syndrome).” (NTSB, 1981, p. 73)*

*“It is a basic precept of psychology and human engineering that the ability of a stimulus to elicit a response (in this case, the ability of a warning tone to get controllers’ attention) is reduced when the stimulus is habitually presented without a reinforcement.2/ Reinforcement for a controller would be the acquisition of useful information from an aural alarm. In other words, when a controller is continually subjected to nuisance alarms, i.e., those that are perceived as useless or distracting, he/she will pay progressively less attention to it.” (NTSB, 1984, p. 5).*

Both of these two passages emphasise the desensitization resulting from over-exposure to nuisance alerts; however, they both lack a clear indication of which system is responsible for generating nuisance alerts. It could in fact be the MSAW, the CA or even another system. A more specific reference to the nuisance-generating alarm can be found in other passages. For instance in a letter issued in 1984, the system held responsible for the majority of nuisance alarm is the CA:

*“According to the supervisor, “a great majority of the time the thing goes off [the CA] and you really don’t pay attention to it 98 percent of the time.” Unfortunately, when a controller ignores the aural alarm, he may be ignoring a warning from the MSAW system, rather than a conflict alert.” (NTSB, 1984, p. 5)*

In this passage, the NTSB essentially attributes the occurrence of nuisance alerts to the CA. Furthermore, it concludes that, as both the CA and MSAW have the same generic aural alarm, desensitisation to the former means desensitization to the latter. In 1990, the problem of nuisance alerts is linked not just to the CA but to the MSAW also:

*“The Safety Board believes that continued exposure to repetitive alerts was at least partially responsible for the delayed action to the conflict alert by the Washington F1 controller...Of all [low] altitude and conflict alerts received in a typical facility, very few actually require intervention by any one controller. As a result, controllers become conditioned by repetitive aural alarms, many of which are not critical.” (NTSB, 1990, p. 5)*

To summarise, the NTSB’s request to differentiate between the CA and MSAW aural tones seems to be justified by the need to enhance the presentation of the aural signal. Such an enhancement would increase the perceptual salience of both alerts. In other words, it



would have captured more effectively way the attention of controllers that were conditioned to ignore alerts due to their frequent exposure to nuisance alerts.

Intrinsic in these considerations is the fact that the NTSB acknowledged the problem of nuisance alerts earlier than 2006—the year in which the first safety recommendation addressing this problem was issued (§ 5.3.2 ). However, before that year the problem is seen as a justification for a design change, i.e., disambiguating the MSAW and the CA aural alerts, rather than as the basis for a safety recommendation mandating the reduction of nuisance alerts.

#### FAA' position: audio alarm serving as attention getter only

The FAA did not concur with the idea that separate aural alarms were needed for the MSAW and CA. All the three times it received this NTSB's request, the FAA stated that the MSAW and CA aural alerts did not need to be differentiated because this would not improve air traffic controllers' response to the alarm. By looking at the passages of FAA's response letters to A-81-134, A-84-83, and A-90-161, the following to FAA's reasons were identified.

First, to the FAA the aural alert warning served as an attention getter only. In other words, the agency saw the alarm signal as serving the sole purpose of directing the controllers' attention towards the radar display—rather than triggering a controller action, i.e., transmitting the alarm to the pilot. This is evident in the following passages:

*"We do not conclude that separate alarms are needed for the low altitude and conflict alerts. We believe that audio alarms represent a general warning or attention getter... The controller does not take control action based on the audio alarm... The alarm or the alarms mean the same thing, scan the display." (FAA letter dated 21, December 1981)*

*"As previously noted, the aural alarm represents a general warning or "attention getter" (statement repeated in FAA letters dated 18<sup>th</sup> January 1985 (FAA, 1985a) and 26 August 1985 (FAA, 1985c).)*

*"The aural alert represents a general warning by identifying an aircraft involved in a potentially unsafe situation, and it requires controllers immediate attention." (FAA letter dated 18 January 1991)*

Second, the FAA justifies the point above—limiting the function of the audio alert to that

of an “attention getter”—on the basis that it is the visual alert that shows the specific nature of the imminent threat, i.e., MAC or CFIT. And, as the visual alert is available on the radar display, it is here that controllers have to look at, after hearing the aural alert, to check the actual nature of the hazard:

*“The blinking alphanumeric represents the specific warning. It identifies the aircraft involved and the nature of the problem.”(FAA letters dated 18<sup>th</sup> Jan 1985)*

*“The blinking alphanumeric represents the specific warning. It identifies the aircraft involved and the nature of the problem, i.e., low altitude or alert or conflict alert. Either situation requires the controller’s immediate attention.” (FAA letters dated 26 August 1985)*

To summarise, the FAA did not see the need to disambiguate the MSAW and CA aural alarms. To the agency the aural alarm served the sole purpose of informing the controllers that an event of interest has occurred that requires immediate verification—on the radar display—of the involved aircraft and of the type and extent of the hazard. No control action can be taken based on the aural alert alone. The FAA, in essence, saw the aural and visual alerts as occurring at two separate stages of a predefined alerting sequence: first, the aural alert attracts the controller’s attention; second, the visual alert informs the controller of the type and the nature of the danger. From this perspective, providing more information by means of the aural alert would have not improved controllers’ response to the alarm as they still had to check the nature of the problem on the radar display.

Also, it can be noted that in maintaining this position the FAA seems not to address the nuisance alert problem: no-mention has been found in the reviewed FAA letters about this problem (see Table 11). The response letters of the FAA appear to address only the adequacy of the NTSB requested change—audio alarm differentiation—, but did not comment on the problem justifying this change.

#### 5.3.3.2. RC7: Amend existing regulations to make mandatory the transmission of MSAW alert to the pilot

In 1984, by means of safety recommendation A-84-84, the NTSB requested the FAA to amend the procedure contained in air traffic control Handbook 7110.65C, paragraph 33 to require that controllers transmit immediately a low altitude alert to any aircraft under their control that has activated the MSAW. This request was issued following the analysis

of a sequence of incidents in which nine aircraft—while approaching Washington National Airport, Washington, D.C—descended below the altitude expected for the approach.

The recommendation letter carrying A-84-84 was followed by an exchange of correspondence letters between the FAA and the NTSB (see Table 13): on the one hand the FAA stated (and restated) its non-implementation rationale for the change demanded by A-84-84; on the other the NTSB restated the importance of implementing it. The respective rationales of the two agencies identified by the analysis are summarized in Table 14 and explained in the next two sections.

**Table 13.** Safety recommendation requesting FAA to modify existing regulations to make MSAW alarm transmission obligatory for controllers. The table reports the related NTSB safety letter and the ensuing letters exchanged between NTSB and FAA.

Safety Recommendation	issued in NTSB safety recommendation letter:	ensuing letters exchanged between FAA and NTSB
A-84-84	A-84-82 through -84, Aug 13, 84	(1) FAA letter Jan 18, 85 (2) NTSB letter July 1, 1985 (3) FAA letter Aug 26, 85 (4) NTSB letter Jan 24, 1986

**Table 14.** Summary of the NTSB’s and FAA’s positions over RC7.

NTSB Rationale	FAA Non-implementation rationale
<p>1-Under existing provisions (Handbook 7110.65D), the controller has the authority to judge the reliability of the alarm;</p> <p>2-However, controllers may (mistakenly) fail to pass a reliable alert to pilots;</p> <p>3-MSAW activation is a definite indicator of unsafe proximity to the ground;</p> <p>4-As soon as the MSAW alert is triggered, controller should not be called to make a judgment to airplane safety, but transmit the alert to the concerned pilot;</p> <p>5-It the pilot who has the responsibility to establish safety, not the controller.</p>	<p>1-Existing procedures were intently designed to let the controller to determine actual aircraft safety prior to issuing the alert;</p> <p>2-When alerted, the controller needs to glance at the visual display prior to delivering the alert;</p> <p>3-To issue an alert in an uncompromising situation could only desensitize pilots to MSAW alerts.</p>

NTSB Position: the transmission of the alert should be mandatory for the air traffic controller

The NTSB rationale for requesting A-84-84 is that controllers should not have the authority to assess the extent of the hazard and decide on whether the alarm had to be transmitted to the pilot—they should pass the alarm to the flying crew as soon as this was generated.

In the text of the concerned safety recommendation letter NTSB observed that existing regulations give controllers the option to judge on the appropriateness of a low altitude alert before transmitting it to the pilots:

*“The FAA’s Air Traffic Control Handbook, 7110.65C, paragraph 33, provides guidance to air traffic controllers to: “issue a safety advisory to an aircraft if you are aware the aircraft is at an altitude which, in your judgment, places it in safe proximity to terrain, obstruction, or other aircraft”. Note 2 in paragraph 33 states in part, “recognition of situations of unsafe proximity may result from MSAW...” However, the Safety Board was given to understand that the phrase, “in your judgment,” gives the controller the option, once the airplane has been identified on the BRITE display, to look at the aircraft from the tower cab and form a judgment concerning the aircraft’s safety. If in the controller’s judgment the airplane is a safe distance from obstructions and terrain, the controller may elect not to issue a low altitude alert.” (NTSB, 1984, p. 6)*

To NTSB, this FAA’s provision could jeopardize the protective function of the MSAW: the controller may fail, in fact, to pass a reliable alert to the aircrew. Thus, to the NTSB controllers should not have the authority to judge on the validity of the alert, but, as soon as an alert is received, they should inform the pilot immediately:

*“The Safety Board is concerned that the provision of paragraph 33 can lead a controller to nullifying the intent and objective of the MSAW system which is to alert a pilot then his airplane at an unsafe altitude.....The NTSB believes [MSAW] activation parameters are definitive indications of unsafe proximity to terrain, and the controller should not be called upon to make a judgment with regard to an*

*airplane's safety...The controller should immediately inform a flight crew of the activation of a low altitude alert, and any decisions concerning the airplane's safety should be made in the cockpit."*  
(NTSB, 1984, p. 7)

The ideas that (i) the authority to judge on MSAW reliability should stay with the pilot and that (ii) the controller should pass the alarm immediately to the concerned crew are further confirmed in the following two NTSB letters. In particular, the second letter confirms that the assessment of the extent of the hazard, and the need to take corrective actions, should stay with the pilot:

*"We note that the FAA disagrees with this recommendation ...[Existing] procedures give the controller the authority to evaluate the extent of the hazard and to pass the warning to the pilot only if he thinks it is necessary. There is a danger of airplanes being inadvertently flown into ground obstructions...Pilots need to be warned when they fly below the minimum safe altitude and controllers, when prompted by the MSAW alerting device, should be required to inform the pilot of the undue proximity to the ground or an obstacle immediately."* (NTSB, 1985)

*"..FAA again has disagreed with this recommendation and is of the view that the procedures contained in the Revised Handbook 7110.65D, Section 1, 2-6, adequately address appropriate controller action in issuing low altitude alerts. These procedures give the controller the authority to evaluate the extent of the hazard and to pass the warning to the pilot only if he thinks it is necessary. We maintain that the controller should be required to advise the pilot of the hazard, and that the pilot should then be responsible for assessing the situation and taking appropriate action."*(NTSB, 1986)

To summarise, NTSB maintained that (1) the controller should not have the authority to judge on the appropriateness of the MSAW alarm; (2) alerts generated by the MSAW are reliable; (3) as soon as generated they should be passed readily to the pilot; and that (4) it is the pilot who has the responsibility to assess the extent of the hazard—and take corrective actions if necessary—and not the controller.

It can be noted that this set of assumptions contrasts in part with the ideas expressed by the NTSB in relation to RC1. In particular the assertion 2—the idea that alerts generated

by the MSAW are reliable—contradicts the assertion—discussed in relation to C1—that the MSAW generated several nuisance alarms. The alarm could not be reliable and unreliable at the same time.

#### FAA Position: authority on the Controller

The FAA did not concur with the request to make mandatory the transmission of the alert to aircrew for controllers. To the FAA controllers should retain the authority to judge whether to pass the alarm to the pilot. This is evident in the following passage, which appeared both in the first FAA response letter to A-84-84 (FAA, 1985b), and the second letter (FAA, 1985c):

*“PRESENT HANDBOOK PROCEDURES WERE DESIGNED TO ALLOW A CONTROLLER TO VISUALLY SCAN TO AN AIRCRAFT’S POSITION AND DETERMINE IF THE AIRCRAFT IS CLEAR OF OR SAFELY AVOIDING TERRAIN OR OBSTACLES. WHEN ALERTED, A CONTROLLER IN THE CAB NEEDS ONLY TO GLANCE AT THE POSITION OF AN AIRCRAFT TO MAKE A DECISION WHETHER OR NOT TO ISSUE AN ALERT. WHEN THE POSITION OF THE AIRCRAFT CANNOT BE DETERMINED VISUALLY, AN ALERT IS BROADCAST TO THE AIRCRAFT. TO ISSUE AN ALERT TO AN AIRCRAFT IN A NONCOMPROMISING SITUATION COULD EFFECTIVELY DILUTE THE URGENCY OF THE MESSAGE.” (FAA, 1985b)(FAA, 1985d)*

Overall, the passage suggests that the FAA’s rationale against A-84-84 can be summarized by the following two points:

- First, controllers’ judgment is needed to assess whether an MSAW alert has to be passed to the aircraft or not to the concerned aircraft. The first three rows of this passage make it clear that the regulation in question is intently designed so to ensure that controllers, when alerted, should first check the specific aircraft on the radar display, and then they should decide whether to transmit the alarm to pilots. In short, controllers should have the authority to decide whether to pass the alert to the pilot.
- Second, doing otherwise—i.e., passing any alert unconditionally to the flying crew—, would risk desensitizing the flying crew to low altitude alerts provides by air traffic controllers. In other words, the FAA seems to imply that controllers should play a filtering role: their judgment is needed to prevent transmitting unreliable alerts to aircrew.

## 5.4. SECOND-LEVEL CODING RESULTS: ORGANISATIONAL PRECURSORS IDENTIFICATION

---

It is now important to elaborate on what these findings mean for the objective of the research, i.e., the development of the theoretical framework of organisational precursors to HAI issues. In particular, the case seems to suggest that the management of a HAI issue appears to depend on the *organisational assumptions driving automation improvement*. By this, it is meant the assumptions held by the organisation about (i) the specific HAI issue and (ii) the role of the automated system in the operational environment—i.e., how the system is to be used by the intended user. These assumptions seems to drive the decisions about how the problems with the system in question can be addressed. The next two sections will justify this conclusion based on the evidence resulting from the first-level coding cycles exposed in the previous section. In particular:

- Section 5.4.1 will summarize the history of the nuisance alert, more specifically, how the nuisance alert problem was viewed by the NTSB over the analysed period;
- Section 5.4.2 will describe the two underling and contrasting views of the alarm role held by the NTSB and FAA.

### 5.4.1. History of the nuisance alert problem in the studied context

The interpretation of even a relatively simple HAI issue like a nuisance alert is not straightforward in complex, high-consequence organisations. Organisations may develop different views about the specific hazard and develop their response accordingly. To support this point it is useful to review the main aspects of the history of the nuisance alert problem that have emerged from the analysis presented in the previous section.

1. *First, the evidence collected in this study shows that the MSAW-related nuisance alert problem could not be analysed as an individual problem, as expected at the beginning of the study:* its trajectory was strictly related to what has been classified as the “controller lack-of-response” concern, i.e., the tendency by controllers to neglect a reliable alert generated by the MSAW. The first, initial analysis (§ 5.3.1) suggested this latter was the main concern that the NTSB noticed regarding the use of the alarm, as this concerned recurred across five safety letters. At the same time the issue of nuisance alerts appeared to be debated as a sub-item in the context of this concern. In particular, it was considered as a contributory cause to this concern together with

issues in the areas of alarm HMI design and response procedures, as shown by the evidence presented under section 5.3.2.

2. *Second, the first NTSB safety recommendation addressing the nuisance alert problem was issued in 2006, after nearly thirty years of operational life of the alarm (§ 5.3.2).*

Up to that year the majority of the NTSB efforts to address the lack-of-response problem focused on requesting changes directed (i) at modifying the MSAW HMI design (10 safety recommendations=77%) and (ii) the procedure defining the response to the alert (2 safety recommendations=15%)—both areas encountering some resistance by the FAA. It is in 2006 that the NTSB acknowledged the (primary) contributory role of the nuisance alert to the lack-of-response problem.

3. *From a safety perspective, this finding is somewhat surprising because it would suggest that the NTSB neglected the nuisance alert problem up to that year.* However, the subsequent analysis of the specific changes contrasted by the FAA (§ 5.3.3) showed this was not the case:

- *NTSB's mentions of the problem were found in 1981, 1984, and 1990; however, in these years the problem was not translated in a corresponding nuisance-alert specific safety recommendation.* It was mentioned, instead, as a justification for another change to the FAA: disambiguating between the MSAW and CA aural alerts (i.e., RC1). This was the case because to the NTSB, this design change would have increased the ability of the aural alerts of both systems to attract the attention of controllers towards the specific risk the alarms were intended to avoid. The reason was that controllers were found to be exposed to nuisance alerts on continual basis, and a more salient aural alert would have been more attention capturing.
- *In turn RC1 received a negative response by the FAA.* To the latter, differentiating the aural alerts would have not enhanced the alerting sequence envisaged for the MSAW (and the CA): both the MSAW and CA aural alerts functioned only as attention grabber; it was the visual alert—displayed on the radar display—that provided information about the type and the entity of the danger. In short, when hearing the aural alert (general alert), controllers had then to check the visual alert (informative alert). To the FAA this envisaged alerting sequence would have not been improved by a more informative aural alert: after hearing it, the controller's check of the visual alert was still needed.



4. *It has to be noted that the NTSB (i) recognition of existence of the nuisance alert problem, noted in relation to RC1 (mentioned here above), contrasts with the recognition of (ii) the MSAW as a reliable system, noted by the agency in relation to RC7.* This latter assumption was used as part of the argument which justified the NTSB request to change existing regulations in order to make more mandatory for controllers the transmission of the alert to the pilot (§ 5.3.3.2). The two positions seems incompatible because the same system, in principle, cannot be both seen as reliable and an unreliable at the same time.

To summarise, up to 2006 two main findings have emerged with the history of the nuisance alert: first, the problem of nuisance alert was known, however not seen as an individual target for improvement by the NTSB, but as a justification for a change—disambiguating between the MSAW and CA aural alert alerts—aimed at solving another problem, the “controllers’ lack of response” concern. This was the main issue with the MSAW (according to the NTSB), not the nuisance alert. In turn, this change turned out to be highly contrasted by the FAA.

Second, the NTSB seemed to maintain a contradicting position about the nuisance alert problem: while recognizing it as a justification for RC1, the agency appeared to neglect it in order to justify RC7. With the available data, it is not possible to explain this contradiction and why and how the NTSB came to realise the importance of treating the nuisance alert problem a standalone problem. However, for the broader objective of the thesis—the development of a framework of organisational precursors to HAI issues—the history of the nuisance alert problem in the US is important because it shows that interpreting the significance of a HAI issue is not necessarily an easy endeavour in safety-critical service provider organisations. While the nature and the significance of the problem may be relatively simple when seen from the perspective of the front-end practitioners that uses the concerned automated system, for stakeholders at higher organisational levels the same problem may turn into an ambiguous (and contrasted) issue—as it happened in the case reported here.

### 5.4.2. Organisational assumptions about the role of the alarm

It can be noted that the behaviours of the NTSB and the FAA analysed in the previous section are grounded on two different views of the alarm role, i.e., views about the way the alarm was intended to function and to be responded by the air traffic controller. The two agencies mandated (and resisted) some changes to the MSAW that were consistent with two diverse underlying views of alarm role. In particular, the NTSB viewed the alarm as a *hazard resolver*; the FAA, as an *attention director*. These two conflicting views are summarized in Table 15, and are described in the following sub-sections.

**Table 15.** Contrasting views regarding the role of the MSAW held by the NTSB and the FAA. References (included in brackets) refer back to the tables of Study 1 in which the particular category was identified.

	NTSB	FAA
<b>Emerging views of the alarm</b>	<u>Hazard resolver</u>	<u>Attention director</u>
<b>Reliability of the MSAW</b>	-The alarm is reliable, it is a definite indicator of an unsafe situation (Table 14, NTSB #3)	-The alarm reliability needs to be verified by the controller (Table 14, FAA #1)
<b>The aural alarm means</b>	-Deliver the alert to pilot (Table 14, NTSB #4) -Avoiding doing so would risk failing to pass a reliable alarm (Table 14, NTSB #2)	-Scan the visual display (to check involved aircraft, type and severity of the danger) (Table 12, FAA #2) -No control action can be taken based on the aural alert alone (Table 12, FAA #3)
<b>Authority</b>	-It is the pilot—not the controller—who has the authority to judge the reliability of the alarm (Table 14, NTSB #5)	-The controller can judge whether the alert has to be passed or not to the pilot (Table 14, FAA #1)

#### 5.4.2.1. NTSB: viewing the MSAW as a “hazard resolver”

The NTSB appeared to view the alarm as a hazard resolver—i.e., an alarm system whose purpose is to direct the controller’s attention towards the implementation of a prescribed response. This statement is grounded on the observation that the NTSB maintained that

upon presentations of the alert signal, there should be no assessment of the alarmed situation by the controller; the controller had to pass the alarm immediately to the pilot. This position finds support in the following evidence:

- First, the analysis of the NTSB rationale behind RC7. That analysis showed that the idea of leaving no authority to the controllers was mentioned explicitly by the NTSB (§ 5.3.3.2);
- Second, the observation of the intent of the two requested changes falling under the procedure area of change (§5.3.2). Both RC6 *Amending regulations to mandate the transmission of MSAW alerts to pilot*, and RC7 *Include supervisors in the alerting loop*, aimed at making more constraining for controllers the NTSB's prescribed response to the alarm. In embedding this intent, both changes were consistent with the idea that no authority should be left to the controller in responding to the MSAW.

#### 5.4.2.2. FAA: viewing the alarm as an “attention director”

The FAA opposed to the NTSB its own view of the MSAW role, i.e., that of an attention director. In other words the FAA viewed the MSAW as a system that directs the controller's attention towards the assessment of a hazardous situation; if the situation is deemed to be hazardous then the response procedure can be activated. In this view of alarm, the controller has the authority to judge on the reliability of the alert. This position has emerged from the following evidence:

- First, by the analysis of the FAA's non-implementation rationale for RC1. This analysis concluded that to the FAA the aural and visual alarm had each a specific function: the former, warning the controller of a general hazard and informing her/him to scan the computer screen; the latter, providing information on the specific hazard (§ 5.3.3.1).
- Second, by the analysis of the FAA's non-implementation rationale for RC7. This analysis concluded that, according to the FAA, the judgment by the controllers on MSAW reliability is necessary in order to avoid passing unreliable alerts to pilots (§ 5.3.3.2).

## 5.5. CHAPTER CONCLUSIONS

---

This chapter has reported the first empirical study of this thesis, which has inquired into the history of the nuisance alert problem related to the MSAW system in the US. As a result of this process the study has identified one category of organisational precursor to HAI issues in complex, safety-critical service provider organisations: *OP1: The organisational assumption driving automation improvement*. This organisational precursor includes the views held by the organisation about the alarm role and the HAI issue itself. In identifying this precursor, the chapter has identified the first category of the emerging OPHAI framework, as shown in bold in Table 16. The same table also anticipates the precursors that will emerge in the next chapter (these are indicated in shaded grey).

*Table 16. The initial version of the OPHAI framework, as resulting from this study.*

<p><b><u>OP1: ORGANISATIONAL ASSUMPTIONS DRIVING AUTOMATION IMPROVEMENT</u></b></p> <ul style="list-style-type: none"> <li>- View of the alarm's role</li> <li>- View of the HAI issue</li> </ul> <p><b><u>OP2: ORGANISATIONAL CAPABILITY FOR HANDLING HAI ISSUES</u></b></p> <ul style="list-style-type: none"> <li>- Parameterisation process</li> <li>- Positive attitude towards air traffic controller's involvement</li> <li>- Supporting software tools and organisational roles</li> </ul> <p><b><u>OP3: CONTROL OVER IMPLEMENTATION QUALITY AT THE BOUNDARY BETWEEN THE SERVICE PROVIDER AND THE SOFTWARE MANUFACTURER</u></b></p>
---

# Chapter 6.

## STUDY 2 Results

### 6.1. CHAPTER INTRODUCTION

---

The case study reported in the previous chapter has provided the first step in the development of the emerging framework of organisational precursors to HAI issues that this thesis aims to deliver. The present chapter continues this development path by shifting the attention to the European context of MSAW implementation. The study consists of a retrospective investigation into the historical trajectory of the MSAW-related nuisance alert problem within the organisational and implementation contexts of two selected air navigation service providers, or ANSPs. The development of this dual-case study is based on interviews, observations and documentary sources. Compared with the previous chapter, the present study focuses more on the organisational dynamics, processes, and decisions internal to the organisation having direct command and control over the MSAW implementation and operation. The chapter is organized as follows:

- Section 6.2 provides on background information useful for the understanding of the case;
- Section 6.3 presents the case-specific findings, i.e., it reports on the categories identified within the two studied organisations, which capture organisational conditions and events relevant for the handling of the nuisance alert problem;
- Section 6.4 links the previous section's findings to the objective of the thesis—the development of a qualitative framework of the organisational precursors to HAI issues. In particular, this section compares these organisation specific findings for the purpose of integrating and subsuming them under the broader categories of organisational precursors these findings stand for. As a result of this process, the section provides an expanded and updated version of the emerging OPHAI framework.

## 6.2. BACKGROUND

---

The MSAW has been introduced in the European Air Traffic Management (ATM) system in the mid-1990s, nearly twenty years after its introduction in the US. As noted in section 4.3.3, in the European ATM system the MSAW is considered as part of a class of applications called ground based safety nets. These are warning systems available to controllers to warn them about an imminent risk to flight safety.

Together with other safety nets, the MSAW has become a safety priority in the EU ATM system following the occurrence of two accidents involving commercial airliners that happened in Europe in the early 2000. The first one involved a mid-air collision between a cargo jet and a civilian passenger aircraft occurred over the Lake Costanza in 2001 (BFU, 2004). The second one, directly relevant for the MSAW, involved a CFIT in 2001, in which a civilian airliner impacted high terrain whilst attempting to perform a missed approach procedure at Zurich airport, in Switzerland (AAIB, 2001). These accidents highlighted problems with the effectiveness of ground based safety nets as implemented and operated by ANSPs at that time. For some of these ANSPs, the MSAW triggered too many alerts resulting from poor parameterisation of the terrain database, which translated into poor acceptance of the tool by air traffic controllers. To address this situation, EUROCONTROL, the European Agency for the Safety of Air Navigation, established a task force concerned with improving the efficacy of MSAW implementations in Europe. This group, called SPIN, launched in 2004 a fact findings reviews across Europe, and since 2005 has promoted the development of specifications of guidance material for MSAW and other safety nets.

It must be noted that, compared to the US, the European context of MSAW implementation is a highly fragmented one. Although required to comply with ICAO international standards and EASA regulations, different European countries have their own ANSP, their own civil aviation authority, and their own accident investigation agency. Furthermore ANSPs can purchase their systems from different software manufacturers. So local state MSAW implementations and safety management practices might differ greatly from country to country. This difference is further exacerbated considering the entrance in Europe of former Russian states since 2000. Overall, this highly fragmented institutional context appears to be a fertile ground in which to investigate organisational precursors to problematic HAI. This chapter investigates the organisational dynamics of two European ANSPs to explore how they handled the “nuisance alert problem” related to their MSAW system.

### 6.3. FIRST-LEVEL CODING RESULTS: ORGANISATION SPECIFIC FINDINGS

This section reports the organisation or case specific findings of this study, which resulted from the first-level coding (see § 4.5.2.1). Such organisation specific findings consist of categories that capture organisational events and decisions that influenced the management of the MSAW and the problem of nuisance alert within Alphasky and Deltasky respectively. Table 17 offers a summary of these organisation specific categories. Alphasky related categories are identified by the letter “A”; Deltasky ones, by the letter “D”. The remainder of this section describes these categories; the next section compares and integrates them in order to extract broader theoretical categories of organisational precursors.

*Table 17. Summary of the categories identified in the Alphasky and Deltasky cases. (For examples of lower level codes and quotations see Table 9.)*

<b>Case 1: Alphasky</b>	<b>Case 2: Deltasky</b>
A1. Having an operational need for introducing the MSAW	D1. Lacking clarity about the operational need for having the MSAW implemented
A2. Clarifying the operational role of the MSAW	D2. Lacking clarity about the operational role of the MSAW
A3 Recognising the problem of Nuisance alert	D3. Purchasing the system in the absence of specified requirements
A4. Parameterisation process;	D4. Manufacturer providing a sub-optimal alarm
A5. Valuing controllers involvement	D5. Realising the nuisance alarm problem after implementation
A6. Supporting software tool and structure of roles	D6. Realising the need for importing parameterisation expertise
A7. Specifying MSAW requirements to the manufacturer	

#### 6.3.1. Alphasky practices

At Alphasky controllers reported to have a positive attitude towards the MSAW. They found it useful in detecting situations in which approaching pilots mistakenly understood to descent to a flight level lower than the one intended and communicated by the controller. Also, the alarm was reported to be useful in periods of high cognitive workload, when managing high levels of traffic. In these situations the MSAW can help the controller

to readily spot on a cluttered radar display an aircraft that is descending too low. The positive attitude of Alphasky controllers towards the MSAW can be explained by looking at the way the nuisance alert problem has been managed within this organisation.

#### A1. Having an operational need for introducing the MSAW

Alphasky had an explicit safety purpose for introducing the MSAW. Senior management decided to adopt the alarm in 1993 as CFIT mitigation measure. The driver for adoption was the need to respond to a safety recommendation issued by the national accident investigation agency (of the country in which Alphasky operates) following the investigation of a CFIT occurred in the early 1990s. In this instance, a civilian airliner, an Airbus A320, crashed into a mountain while performing the approach descent to its destination airport under low visibility conditions. This accident was the last in a series of three major aviation disasters involving commercial aviation aircraft impacting terrain during the final approach phase. The safety recommendation in question demanded that:

*“...a particular effort [has] to be made as soon as possible [by Alphasky] to complete the design and implementation by air traffic control services of a ground based system for detection of an aircraft in dangerous proximity to terrain, whenever technically possible.”*

From this point onwards, Alphasky senior management regarded the introduction of the MSAW as a high priority safety action. The initial plan was to have the system operational in approach control centres by the end of 1995.

#### A2. Clarifying the operational role of the MSAW

At the time the decision to introduce the MSAW was taken, the alarm was not novel to the organisation. Alphasky had initiated the development of MSAW operational requirements in 1988, based on a survey of operational personnel and consultation with accident investigators. Based on the consideration of the US experience with the alarm, it was noted that MSAW could be implemented either as either (i) a hazard resolver, with the MSAW signal transmitted readily to the pilot by the controller, without intermediate controller assessment of the severity of the situation and the reliability of the alarm; or as (ii) an attention director, with the warning signal to be assessed by the controller prior to being transmitted to the pilot. Eventually Alphasky opted to implement the system as a hazard resolver. This is reflected in the following passages taken from an Alphasky's documentary source:

*“[Alphasky] has chosen to develop a reliable system in which all*



*alerts are justified, and which should entail no situational analysis work by the controller, but rather a reflex action informing the aircraft concerned...” (Alphasky documentary source)*

In establishing this goal, also the organisation acknowledged that:

*“The appearance of a large number of false alerts and undesirable alerts can be a factor in a certain dilution of controller vigilance and reduced confidence in the system, which in the end means that the controller-system structure no longer correctly provides the collision avoidance alert service.” (Alphasky documentary source)*

### A3. Recognizing the problem of nuisance Alert prior to implementation

These passages reflect Alphasky’s intent to implement a MSAW alert that was reliable from an air traffic controller perspective, so that s/he could transmit it readily to the flying crew without any intermediate situation assessment. Alphasky seemed to recognise the threat posed by frequent nuisance alerts to the MSAW protective potential: if exposed to this condition controllers would have lost confidence on the MSAW signal. Such a view over the MSAW role lives in the organisational measures established within Alphasky to control and mitigate the nuisance alert problem.

### A4. Parameterisation process

The presence of a formalized and repeatable MSAW parameterisation process is arguably the most visible organisational response developed by Alphasky against the MSAW nuisance alarm problem. Such a process is repeated every time MSAW implementation at a new site begins, and falls within the responsibility of the Alphasky internal Design and Validation team.

This parameterisation process entails three main phases. In the first phase, a sample of live air traffic data from the selected site is recorded, and subsequently this is fed into an initial MSAW configuration. This first configuration usually generates a high number of nuisance alerts both for the final approach path in the proximity of the airport and in the terminal area. So, the MSAW engineer can play with MSAW algorithm and inhibit areas to maximise the warning time while reducing the nuisance alert rate. The use of a sample of recorded real traffic data, instead of automatically generated traffic, is deemed very important in ensuring tuning reliability—i.e., minimising the risk of missing relevant conflict points while reducing the rate of nuisance alerts.

After a first MSAW configuration is defined, the MSAW is installed on the local site. This

system is then fed with real time air traffic data although it is not yet broadcast to controllers. This phase allows for further refinement of the initial MSAW parameters and inhibition area. Such refinements were not possible using the recorded data alone. This phase terminates when the rate of false alerts reaches the threshold of one-per-day.

The third stage consists of operational on-site testing. Here, the MSAW is broadcast to air traffic controllers; however, the alert is not transmitted to the pilots until the system performance is considered acceptable from the perspective of the controllers involved in parameterisation.

### A5. Positive attitude towards air traffic controllers' involvement

The process makes use of input from selected controllers permanently allocated to the safety net domain, who have several years of operational experience combined with knowledge of the technical decisions involved in parameterizing the MSAW. Their judgment is needed to advise engineers about which MSAW alerts are relevant from an end-user perspective.

In a typical engineer-expert controller consultation, the engineer and the controller sit together in front of a chart displaying the final approach segment, as well as the MSAW alerts generated by the aircraft tracks that have descended along this path over a given period of time. Some alerts might be on, some above, some below such path. The expert controller provides his/her best guess about which of these alerts are important from an operational perspective. This allows the engineer to classify which alert is (i) necessary, i.e., "the situation involves a serious deviation below the safety altitude"; is (ii) desirable, i.e., "although there was not a serious deviation below safe altitude, an alert would have been useful in drawing the attention of the controller to the situation"; or is (iii) unnecessary, i.e., "the alert is a genuine nuisance". In turn, the engineer can further fine-tune the shape of inhibition areas, the gradient of glide slopes, and the setup of MSAW parameters to ensure that unnecessary alerts do not get displayed to pilots.

It can be noted that the frequency of these consultations and the length of the implementation process are not based on a fixed formula but depend on the terrain topography and the traffic volumes specific to the site where a new system has to be implemented. The more complex the terrain topography, the more the demand placed on the set up of the grid; the smaller the traffic flow, the longer the time needed to collect a statistically representative traffic sample of approaching aircraft. (For instance, these two conditions are typical of small airports located in mountainous regions, for which parameterisation might take up to two years.)

### A6. Supporting software tools and organisational roles

Since the decision to adopt the system was made, one of Alphasky priorities was to have in place a demonstrator able to visualize the alert on a radar display and tools for recording, flight path analysis and statistics. Such a demonstrator was installed at Alphasky R&D facilities in 1993, and in the first implementation site in 1995, to provide an infrastructure for checking and improving MSAW performances. By means of the demonstrator MSAW alerts could be checked for aspects such as their reliability, efficiency, and operational acceptability.

In addition to the demonstrator, other supporting software tools are available at local site level to monitor/improve the MSAW performances. A software tool has been built that allows refinement of the MSAW inhibition areas. This tool allows the safety net engineer to specify the geographic coordinates and the height of the areas in proximity of the airport in which the MSAW alerts have to be inhibited. Such coordinates are then passed to the Validation and Design team that can approve and implement the required change. This passage is needed as local airport engineers do not have the authority to modify the operational MSAW. Other software for retrospective safety analysis has been found within the quality of service department. This application has a database of site-specific incidents, whenever an incident is selected, it shows a top and a side view of the traffic situation being analysed complete with recordings of the conversations between pilots and controllers. Although not specifically built for the MSAW, this application has been modified to permit the replay of MSAW warnings and other safety nets. This is particularly helpful to investigate those events in which aircraft lost separation from terrain, so as to understand why it happened, who were the responsible actors involved, and what was the role of the MSAW in that particular incident.

According to the interviewed staff, organisational arrangements at national level ensure a focus on constant MSAW monitoring and improvement. In addition to the Design and Validation team, and the controllers involved in parameterisation, Alphasky has in fact created the role of a national safety net leader who ensures visibility of the safety nets issues to senior management and ensures that appropriate funding and staffing are allocated to MSAW and other safety nets.

### A7. Specifying requirements to the manufacturer

The development and implementation effort of Alphasky MSAW involved a joint effort between Alphasky and its national manufacturer. Since the decision to adopt the system was taken, this manufacturer has supplied the earlier demonstrator and other tools

needed to parameterize the experimental MSAW. Also, the manufacturer has developed the operational MSAW based on the requirements specified by Alphasky's Design and Validation team. Today, this manufacturer improves the system regularly based on the requirements specified by the Design and Validation team. These are based on the issues with the alarm and other safety nets collected from the different air traffic control centres in which the alarm is implemented. Having described the practices through which Alphasky handles the problem of the MSAW-related nuisance alerts, the next section will move to the description of the conditions found within the second organisation studied, Deltasky.

### **6.3.2. Deltasky practices**

At the time of the site visits Deltasky was found to have no MSAW system installed. This ANSP implemented the alarm in 2004 and had to remove it due to the excessive rate of nuisance alerts it triggered. Controllers reported that the alarm used to go off for almost every single aircraft engaged in the final approach phase; therefore, the alarm was interfering with their ongoing practices and they requested senior management to turn it off. Within one month of operation, the MSAW was removed from operations. The following narrative reports on the organisational decisions and conditions that, while accompanying the first MSAW implementation in Deltasky, also seem to explain how the organisation inadvertently induced the nuisance alert problem.

#### **D1. Lacking clarity about the operational need for introducing the MSAW**

Deltasky acquired and installed the MSAW as a commercial-off-the-shelf (COTS) system in 2004. The driver for adoption was modernization, although not of the MSAW specifically but of the major software infrastructure and physical facilities. This ANSP operated legacy systems nearly ten years older compared to industry standards. The main problem was the lack of reliability of aging radar and flight data processing systems that caused breakage and interruptions of air traffic service on a nearly daily basis. As stated by the project manager managing the modernization programme:

*“Our burning issue was to have a stable [radar processing and flight data] system. The old one was unstable, dying in our hands virtually... So, our burning issue was to have stable system based on a new pc platform and flat displays, good flight plan processing... We did not have any external driver for implementing MSAW and other safety nets; we needed to catch up with industry standards.” (Deltasky*

*project manager)*

Because the replacement of the radar and flight data processing systems was the “burning issue”, Deltasky focused on specifying their requirements, selecting a software manufacturer that could supply them. At the same time, other top priority areas of change running in parallel included the refurbishing of the approach and en-route control room facilities, and the civil engineering work needed to build and refurbish a new control tower. These were the priorities within Deltasky at the time the MSAW was adopted, rather than the adoption of the MSAW or other safety nets per se. This is reflected in the following passage:

*“The main focus was on the implementation of the new systems to get a higher level of flight data processing capability. We knew pretty well the level of automation that was needed. On the other hand, safety nets were a side dish, we had very little knowledge of them, and they were just a little feature that was offered by \*\*\*[manufacturer’s name omitted] as a part of a larger system. Also, we were not aware of the necessity of having safety nets. This is reflected in the way we parameterised safety nets later on.” (Deltasky project manager)*

## D2. Lacking clarity about the operational role of the MSAW

In addition to highlighting the areas receiving most of Deltasky attention, the third last row of the previous transcript (“Also we were not aware of the necessity of having safety nets”) captures the limited knowledge of safety net purpose and the need for having them implemented within the Deltasky at the time of MSAW adoption. This is indeed supported by other passages from the interviews conducted:

*“They were just new kids on the block, my knowledge of the MSAW was limited to what was specified in ICAO standards. It is through the participation in SPIN we learned that the sole purpose of MSAW was CFIT protection.” (Deltasky project manager)*

*“We saw the MSAW just as a minor technical system, something like, something that controllers would have barely noticed. We did not imagine the mess lying ahead with all the nuisance alerts and the need to parameterize the system.” (Deltasky engineer)*

The expressions our participants used to describe the MSAW (and other safety nets)—i.e.,

“side dish”, “little feature”, “new kids on the block”, “minor technical system”—render (i) the perceived uncertainty about the purpose of this system; (ii) the belief that this was nearly “transparent” to the air traffic controllers, i.e., its introduction brought no impact on the practices of the end-user; and (iii) the belief that it was a relatively simple system if compared to the larger software systems that were being purchased.

#### D3. Purchasing the system in the absence of specified requirements

The above mind-set is evident in the decision to adopt the MSAW, which was taken after the software manufacturer proposed to include the alarm as an added feature to the larger software package under acquisition. While accepting this system, it was reported that only the acronym of “MSAW” appeared in the contract signed with the manufacturer—no specification of system and parameterisation requirements were provided. In other words, the system was purchased as a commodity or black box. The underlying assumption was that the manufacturer possessed the competence needed to implement the alarm autonomously, i.e., with no input from the ANSP side. The manufacturer had been selected among one of the best in the industry, in fact, and its ability to implement a system of relatively low complexity was not questioned.

#### D4. Manufacturer providing a sub-optimal alarm

In retrospect, it emerged that neither the ANSP nor the manufacturer was fully aware of the trade-offs involved in setting up MSAW parameters and the terrain database. The manufacturer re-used a terrain database that Deltasky requested for another (weather-related) function. This latter function was well-known within Deltasky, so that, at the time of purchasing, both the needed HMI features and, in particular, the terrain database were specified for this function. In turn, the manufacturer developed a digital version of this grid into the radar system based to the specifications provided by Deltasky; however, it also proposed to re-use the same grid also for the MSAW, assuming that this was appropriate. Eventually, while the specified grid served well its intended function, its units were far too large for the alarm:

*“The engineers of \*\*\* [manufacturer’s name omitted] deemed the [specified] grid to be okay. However, when we started using it we realised it was so unrefined. If you cannot refine your warning altitude to distinguish between an aircraft that is on a correct course of approach from one that is not, then you have a nuisance.” (Deltasky IT engineer)*

#### D5. Realising the nuisance alert problem after implementation

With no rationale and no requirements specified by the service provider for MSAW parameterisation, and the parameterisation process left to the manufacturer, the quality of the alarm appeared to reflect more the software vendor view than that of air traffic controller's operational needs. When the system went operational it generated a high rate of nuisance alerts that most controllers considered disruptive. This was also the stage at which management realised the problem:

*"It was when the manufacturer had packed its suitcase and gone that we realised we did not have access to system functions, the knowledge, the tools, and the competence to do the remaining implementation process on our own." (Deltasky project manager)*

#### D6. Realising the need for importing parameterisation expertise

The passage above highlights the difficulty perceived within Deltasky about improving its MSAW after the "departure" of the manufacturer. After the manufacturer had fulfilled its contractual obligations, the ANSP did not have an in-house specific team or people with experience in MSAW parameterisation to improve the system. During the implementation, the ANSP personnel were not trained regarding the functionalities and the parameters of the implemented system as this was not part of the contract. Hence, internal ANSP's engineers were just at the beginning of a steep learning curve. Facing these difficulties, Deltasky removed the alarm from operation after its first troubled implementation, and, at the time of data collection, was collaborating with another software manufacturer in order to implement a new set of safety nets.

#### **6.4. SECOND-LEVEL CODING RESULTS: ORGANISATIONAL PRECURSORS IDENTIFICATION**

---

The previous section has reported the organisational experiences of implementing and managing the MSAW systems found in two European ANSPs: a successful case, Alphasky, and a less successful one, Deltasky. That section identified the components of these experiences, i.e., organisational events, dynamics, and interactions associated with the best-in-class and problematic management of the nuisance alert problem. The present section reports on the results of the cross-case analysis of these components (see also § 4.5.2.2). This analysis has explored if and how the categories identified could be subsumed under broader theoretical categories of organisational precursors to HAI issues. As a result of this analytical step the following theoretical categories of organisational precursors have emerged:

- OP1: Organisational assumptions driving implementation and improvement;
- OP2: Organisational capabilities for handling HAI issues;
- OP3: Control over implementation at the boundary between the service provider and the software manufacturer.

The matrix in Table 18 shows how the case-specific categories connect to these three theoretical categories. These categories are presented in the remainder of this section.



**Table 18.** Organisational precursors to HAI issues resulting from the cross-case analysis.

	<b>Alphasky categories</b>	<b>Deltasky categories</b>
<b>OP1. Organisational assumptions driving automation implementation and improvement</b>		
	A1. Having an explicit operational need for introducing the MSAW; A2. Clarifying the operational role of the MSAW; A3. Recognising the role of the nuisance alarm problem prior to implementation;	D1. Lacking clarity about the operational need for introducing the MSAW; D2. Lacking clarity about the operational role of the MSAW; D5. Realising the nuisance alert problem after implementation;
<b>OP2. Organisational capability for handling nuisance alarms</b>		
	A4. Parameterisation process; A5. Positive attitude towards controllers' involvement; A6. Supporting software tools & organisational roles;	D6. Realising the need for importing parameterisation expertise;
<b>OP3. Control over implementation quality at the boundary between the service provider and the software manufacturer</b>		
	A7. Specifying MSAW requirements to the manufacturer.	D3. Purchasing the system in the absence of specified MSAW requirements.

### 6.4.1. OP1. Organisational assumptions driving implementation and improvement

This theoretical category refers to the assumptions that organisational members hold with regard the automated system to be implemented, of how the system has to be used, and the range of potential accompanying HAI issues. In an ideal situation the organisation mobilizes resources to the clarification and specification of the role of the automated tool to be introduced, of why it is needed, of how it will be used by the human operator, of what are the (main) HAI issues. Specifically to Alphasky, three categories have emerged from the first-level coding in support of this precursor:

- A1. Having an explicit operational need motivating the introduction of the MSAW;
- A2. Clarifying the role of what an MSAW system is;
- A3. Acknowledging the problem of the nuisance alerts.

On the negative side, an organisation may commit to adoption of an automated system without exploring and defining the role of the system. In other words it may treat it as a relatively unproblematic “black box”. This mind set can blind the organisation to the intricacies of automation implementation; in particular it can blind the organisation to the range of undesirable HAI issues associated to the introduction of the automated system. In turn, these may become evident after the system has entered into service. Specifically to Deltasky the following dynamics have emerged in support of this pattern:

- D1. Lacking an operational need for adopting the MSAW;
- D2. Lacking clarity about the operational role of the MSAW;
- D5. Realising the nuisance alert problem after implementation.

Note that the definition of organisational assumptions formulated here relates to and expands the results of the previous chapter. Also in that case (the US MSAW) it was concluded that the organisational assumptions about the role of the automated alarm and the accompanying HAI issues seem a relevant organisational precursors to HAI issues, as it was the NTSB’s and the FAA’s organisational assumptions that drove the responses of the two agencies to the problems identified with the MSAW. Thus both studies appear to converge on the idea that the human-centeredness of technology operated by practitioners at the sharp-end may depend on the views that senior management and other stakeholders at the blunt end hold about the technology. This behind-the-scene view can influence the (lack of) mitigation of a HAI issue both before the automated system is introduced in the operational environment—i.e., at the time of implementation

(Study 2)—and after system introduction—i.e., during the post-implementation phase (Study 1).

#### **6.4.2. OP2. Organisational capability for handling HAI issues**

The handling of an HAI issue seems dependent on the availability of a specific organisational capability within the organisation introducing the automated system. By “organisational capability” it is meant the ensemble of the practices—i.e., what the organisation does—that allows the organisation to cope with the demands introduced by automation. The view of capabilities defined here builds on the notion of organisational capability found in the innovation and business literature. For instance, to Leonard-Barton, an organisational capability includes “the system of activities, physical systems, skills, and knowledge bases, managerial systems of education and reward, and values that create a special advantage for a company or line of business.” (Leonard-Barton, 1998). To Prahalad and Hamel, a capability encompasses the “bundle of skills and technologies that enable the company to provide a particular benefit to customers” (Prahalad & Hamel, 1990). To Leinwand and Mainardi, a capability identifies the organisation ability to “reliably and consistently deliver a predefined output, something that “is ensured through the combination of processes, tools, knowledge, skills, and organisation, all focused on meeting the desired result” (Leinwand & Mainardi, 2010). None of these definitions of organisational capability are defined in relation to HAI management in complex, high risk organisations. Regardless of their differences, these definitions capture the intrinsic systemic essence of an organisational a capability: a capability is something that the organisation does over time by means of distinct, although interrelated, components.

In this specific study, the components of such a capability include the availability of a repeatable parameterisation process, of tools supporting this processes, of deference toward the involvement of the end-user, and of a dedicated supporting structure of roles. Such components seem important in order to ensure that a HAI issue receives appropriate mitigation and control actions by the concerned organisation. The case of Alphasky provides evidence in support of this idea. Consistently with its defined view of the alarm, Alphasky had directed resources towards the implementation of that view and the mitigation of the main threat to it, i.e., the frequent generation of nuisance alerts. In particular the following categories were identified within Alphasky:

- A4. Parameterisation process;
- A5. Positive attitude towards air traffic controllers’ involvement;
- A6. Supporting software tools and organisational roles.

On the other hand, Deltasky was not found to possess such a handling capability when it adopted the MSAW alarm. Due to its lack of clarity about the role of MSAW in operation, this organisation appeared to lack sensitivity to the problem of parameterizing the alarm to its own terrain conditions at the time of adoption. It is only when the alarm became operational that senior management realised the need to import capabilities that were available outside of the organisation: during the purchasing and implementation of the system, no specific training was provided to Deltasky engineers for parameterizing the system (see *D6. Realising the need for importing parameterisation expertise*). To summarize, the findings of Study 2 suggest that the understanding of a specific HAI issue appears to require the consideration of the organisational capability which have been established for managing this issue.

#### **6.4.3. OP3. Control over implementation quality at the boundary between the service provider and the software manufacturer**

The two theoretical categories of organisational precursors previously identified, OP1 and OP2, point, essentially, at organisational dynamics that are internal to the organisation. The third category presented here, instead, points at the unfolding relationship between the software manufacturer and the organisation at the time the system is purchased and implemented. From the collected evidence, it emerged that neither Alphasky nor Deltasky developed the system in-house, but contracted out its implementation. They appeared to differ, however, in the way they managed the quality of the automated alarm implementation at the boundary between the service provider and the software manufacturer. By control over the implementation quality it is meant the type of mechanisms that the service provider adopts to ensure that the quality of the supplied system, as provided by the software manufacturer, is fit for its own, specific operational environment (Kern & Willcocks, 2001). As suggested by research on information system outsourcing, such controls may come in a variety of forms—e.g., project plans, contract forms, peer pressure, etc. (Kirsch, Sambamurthy, Ko, & Purvis, 2002)—and establishing appropriate controls is one of the main challenge for the purchasing organisations. In the present case two basic types of controls have emerged, which seems to influence the quality of the alarm implementation.

- **Requirement based.** On the one hand a service provider can recur to the formal specification of the outcome, i.e., system requirements. The analysis of the Alphasky case, the best in class implementer, indicated that this service provider has an internal team responsible for the collection, the specification, and the transmission of the requirements to the software manufacturer (see *A7*.

*Specifying system requirements to the manufacturer*). By means of the specified requirements, the organisation is able to effectively being provided with the system they specify for their operational system;

- **Trust based.** On the other hand, failing to see the complexity of system implementation while trusting the software manufacturer may result in leaving the quality of implementation to the latter. After deciding to implement the MSAW after this was proposed by the manufacturer, Deltasky did not specify the requirements of the alarm to the manufacturer in the initial contract, as the manufacturer was trusted because of its reputation (see *D3. Purchasing the system in absence of specified MSAW requirements*). However, it is generally acknowledged that in the absence of adequate outcome controls such as system requirements, other means need to be in place to ensure that the manufacturer, the controlee, will act in the best interest of the controller, the service provider. Failing to do so risks leaving the organisation purchasing the system with an implementation reflecting more the manufacturer view of implementation than the client side. The case of Deltasky seems to confirm this general pattern: the absence of formal controls over the alarm implementation at the service provider-software manufacture boundary resulted, in fact, on the provision of a suboptimal system—in that specific case a system with a terrain grid specified for another function (see case category *D4. Manufacturer providing a sub-optimal alarm*).

Overall, the difference observed between Alphasky and Deltasky regarding their interaction with their respective manufacturers, and, in particular, the problematic outcome of the trust based approach observed in the Deltasky case suggest that the control over implementation quality at the service provider-software manufacturer boundary is an organisational dynamic that one has to consider to understand the quality of human automation interaction of an operational system. The broader implication is not that all trust based form of control necessarily result in the poor implementation of a supplied safety-critical alarm system. But, indeed, the collected evidence points at the potential risk associated with this approach—this approach may inadvertently result in a suboptimal implementation quality, something that is not realised by the purchasing company until the operational system is implemented and when the manufacturer has completed its contractual duties.

## 6.5. CHAPTER CONCLUSIONS

This chapter has presented the second empirical study of this thesis. The study examined the organisational experiences of (mis) managing the issue of the MSAW-related nuisance alerts during the implementation and operation of the MSAW within two European ANSPs. In particular the chapter has looked at how these two organisations have handled the problem of the MSAW-related nuisance alerts. The chapter has contributed to OPHAI development by:

- Further supporting the importance of *OP1: Organisational assumptions driving automation implementation and improvement* as a precursor to the management of HAI issues. This theoretical category was already identified at the end of Study 1 (chapter 5), and it has received further support by the evidence collected in this study;
- Identifying two additional relevant theoretical categories of organisational precursors, namely: *OP2: Organisational capability for handling HAI issues*, and *OP3: Control over implementation quality at the boundary between the service provider and the software manufacturer*.

The updated version of the OPHAI framework resulting from this study is presented in Table 19. What remains to be done is to assess the validity and generalisability of the framework. This is undertaken in the following chapter.

**Table 19.** The updated version of OPHAI based on the findings of Study 2. While OP1 has remained unchanged, two more categories of precursors (OP2 and OP3) have been identified.

<p><b><u>OP1: ORGANISATIONAL ASSUMPTIONS DRIVING AUTOMATION IMPLEMENTATION AND IMPROVEMENT</u></b></p> <ul style="list-style-type: none"> <li>- View of the alarm's role</li> <li>- View of the HAI issue</li> </ul> <p><b><u>OP2: ORGANISATIONAL CAPABILITY FOR HANDLING HAI ISSUES</u></b></p> <ul style="list-style-type: none"> <li>- Parameterisation process</li> <li>- Positive attitude towards air traffic controllers' involvement</li> <li>- Supporting software tools and organisational roles</li> </ul> <p><b><u>OP3: CONTROL OF IMPLEMENTATION QUALITY AT THE BOUNDARY BETWEEN THE SERVICE PROVIDER AND THE SOFTWARE MANUFACTURER</u></b></p>
--

# Chapter 7.

## **STUDY 3 Results**

### **7.1. CHAPTER INTRODUCTION**

---

The present chapter reports the results of the third and final empirical study of the thesis: a subject matter expert (SME) study aimed at refining and expanding the initial version of the OPHAI framework, which emerged from the previous two studies. The study subjected the definitions of the theoretical categories of precursors identified to the criticism of a sample of 11 subject matter experts (SMEs), defined as individuals with professional experience in the implementation and management of the MSAW and other safety nets.

The study aimed to explore the SMEs' reactions to the framework. It was anticipated that, if the categories of organisational precursors identified resonated with real life situations experienced by these experts, they would prompt triggered comments and criticism, which in turn would have clarified, refined and expanded the categories of the framework.

The chapter is organised as follows:

- Section 7.2 presents the results of the study;
- Section 7.3 discusses the implications of the findings of this study for the emerging framework.

## 7.2. RESULTS

The SMEs' feedback was collected by means of a group discussion and an individual questionnaire (see § 4.4.3). At the start of the group discussion, the participating experts were given an introductory presentation about the framework and its categories of precursors, as resulting from Studies 1 and 2. At the end of the presentation they were administered an individual questionnaire, which they were asked to return via e-mail after the meeting. These data were then coded to identify corroboratory categories, i.e., data segments that captured relevant ideas, opinions, attitudes, and analogous situations mentioned by the experts. This process (§ 4.5.3) led to the identification of 11 corroboratory categories, as reported in Table 20. These are described in the following three sub-sections.

**Table 20.** Study 3: list of the corroboratory categories emerged in relation to the three categories of organisational precursors of the OPHAll framework.

<b>CORROBORATORY CATEGORIES RELATIVE TO OP1</b>
<ul style="list-style-type: none"> <li>- <b>OP1.1: Senior management tendency to overlook safety nets intricacies</b></li> <li>- <b>OP1.2: Missing conceptual link between safety and safety nets</b></li> <li>- <b>OP1.3: Safety nets used for purposes other than the intended one</b></li> <li>- <b>OP1.4: Developing a correct understanding of safety nets</b></li> </ul>
<b>CORROBORATORY CATEGORIES RELATIVE TO OP2</b>
<ul style="list-style-type: none"> <li>- <b>OP2.1: Positive attitude towards air traffic controllers' involvement</b> <ul style="list-style-type: none"> <li>o Establishing mechanisms for collecting controllers' feedback</li> <li>o Failing to establish adequate mechanisms</li> <li>o Sources of lack of reporting               <ul style="list-style-type: none"> <li>- Report not triggering corrective actions</li> <li>- Inhibiting command and control culture</li> <li>- Blaming controllers for safety nets problem</li> </ul> </li> </ul> </li> <li>- <b>OP2.2: Senior management support to the safety net domain</b> <ul style="list-style-type: none"> <li>o Important to ensure long term sustained organizational focus over the domain of safety nets and the accompanying problems</li> <li>o Promoting the development and retention of expertise</li> <li>o Promoting awareness of safety net's role within the organisation</li> </ul> </li> <li>- <b>OP2.3: Safety nets team</b> <ul style="list-style-type: none"> <li>o Safety nets team components</li> <li>o Essential to ensure coordination</li> <li>o Small ANSPs lacking the resources to have a dedicated safety nets team</li> </ul> </li> <li>- <b>OP2.4: Safety nets specific performance framework</b> <ul style="list-style-type: none"> <li>o Useful for monitoring data</li> <li>o Important for defining reference values</li> <li>o Needs to be safety nets specific</li> </ul> </li> </ul>
<b>CORROBORATORY CATEGORIES RELATIVE TO OP3</b>
<ul style="list-style-type: none"> <li>- <b>OP3.1: Patterns common to other less successful implementation contexts</b></li> <li>- <b>OP3.2: Patterns affecting also large ANSPs</b></li> <li>- <b>OP3.3: Working closely with the manufacturer to avoid the problem</b> <ul style="list-style-type: none"> <li>o Formal requirements specification</li> <li>o Difficult for ANSPs to define safety nets requirements in absence of expertise</li> <li>o Framing implementation in the context of a continuous relationship</li> </ul> </li> </ul>



The three subsections follow the same structure in order to make it clear how the corroboratory categories impact each organisational precursor. Each section provides:

1. A brief description of the initial definition of the organisational precursor, as resulting from the previous chapter;
2. The presentation of the actual results of Study 3, namely the corroboratory categories that have emerged from the study;
3. A discussion of the implications the corroboratory categories have for the organisational precursor.

### **7.2.1. SME feedback on OP1: Organisational Assumptions Driving Automation Implementation And Improvement**

This precursor captures the fact that the management of HAI issues may depend on the assumptions that service provider personnel holds about the role of the alert and purpose, and the HAI issue itself (§ 6.4.1). The way these organisational members think about safety nets is important because it can shape the way they conceive and implement them. The following list reports the corroboratory categories emerged in relation to OP1: these capture the reactions of the SMEs that have been triggered by this initial definition of the framework.

1. **Senior management tendency to overlook safety nets intricacies.** The SMEs confirmed to have experienced situations in which safety nets have been regarded by senior management and personal engaged in their implementation as relatively simple systems, essentially IT systems which require software expertise only for their implementation (SME3; SME7).

*“Also today, safety nets, despite their names, tend to be viewed as minor systems, as something that has to be dealt with by the IT engineer” (SME7).*

The consequence of this thinking is that the definitions and meaning of the system may not be elaborated by developers and management prior to its implementation (SME4; SME6). The organisation, in turn, may remain blind to the range of issues that the alarm may bring in the field of practice.

2. **Missing conceptual link between safety and safety nets.** Further illustrating the vagueness that may accompany the introduction of safety nets, three experts (who have been involved in past safety nets implementation in five ANSPs) remarked that although the label of “safety nets” would suggest an intrinsic and very strong

conceptual link with safety, for some providers this link was not necessarily established, e.g.:

*“From our experience safety nets were unfortunately not that important in the past... one of the problem is that it was not even clear the relationship between safety nets and safety, (and also today this relationship is not [always] clear).” (SME4)*

*“If you call a system ‘safety net’ it is normal to expect this system to bring some improvement in the area of safety; however, it is not uncommon to fail to associate safety nets to safety for ANSPs not familiar with these systems.” (SME6)*

Given the available data is not possible to further clarify the way in which this conceptual link can be (mis) construed by practitioners. Arguably, this may reflect the fact that service providers have as a mission the provision of air traffic services. The implementation of new systems is something that adds on the top of this mission; therefore, it is reasonable to expect that the organisation allocates relatively limited resources to exploring the meaning that the specific system may have in the organisational context.

3. **Safety nets used for purposes other than the intended one.** To further stress the vagueness about safety nets role and purpose, it was reported that safety nets, in some service providers, were used for purposes different than the officially intended one. Two providers, for instance, were reported to have used the STCA not as a safety net but as a control aid or decision support tool: controllers used the alarm to check whether a possible change in flight level would have triggered an alert (SME2; SME5; SME7), when in fact safety nets should be used as a warning only—i.e., warning that some predefined threshold conditions have been exceeded. A third provider was reported to have claimed on its web site to have achieved a gain in traffic capacity after having installed safety nets, despite such a claim contradicted the intended purpose of safety nets, which is risk avoidance, and not production increase.

It was reported that situations such as these were observed especially when safety nets were first introduced in the European Air Traffic Management system, in the early 90s. At that time, their role was only minimally defined by international standards, and this contributed to the relative uncertainty about their purpose and role (SME2; SME4; SME10). SMEs’ widespread conviction was that this situation began to improve in 2005, when EUROCONTROL established the SPIN network (already mentioned in § 6.2).

4. **Developing a correct understanding of safety nets.** The SMEs supported the idea that

senior management understanding of safety nets role is important to ensure that resources are activated and directed towards the safety nets domain. In general, best in class implementers allocate an initial time period to explore and clarify the purpose of the safety net to be introduced. This implies conducting, as a minimum, activities such as gaining understanding of how these systems have been implemented elsewhere, specifying the requirements of the systems, and conducting a cost benefit analysis. One expert noted that some ANSPs—which went through this process and eventually come to the conclusion that the safety nets were not needed in their operational context—may have a clearer understanding of the purposes of safety nets than other ANSPs that have adopted these systems without familiarising first with them (SME2).

### 7.2.1.1. Implications for OP1

The initial definition of OP1 appeared to resonate with the professional experience of the SMEs involved in this study. This initial definition evoked, in fact, analogous professional situations the experts have experienced in the past, in which they have observed this precursor at play. The SMEs confirmed to have experienced (1) situations in which senior management committed to implement safety nets without looking into their intricacies; (2) situations in which the conceptual link between safety and safety nets was missing; and (3) situations in which safety nets were used for purposes other than the intended one. Also, they confirmed that (4) the development of a correct understanding of the role of safety nets is an important part in the introduction of these systems. In conclusion, the feedback collected in relation to OP1 tends to reinforce the confidence on the definition of this precursor, as formulated at the end of the previous chapter. Overall, this feedback adds corroboratory support to the idea that the way organisational personnel at the blunt end think about automation influences the way automation and the accompanying HAI issues are managed.

*Table 21. Evolution of OP1 as a result of Study 3: this category and its sub-items have remained unchanged.*

<p><b>OP1: ORGANISATIONAL ASSUMPTIONS DRIVING AUTOMATION IMPLEMENTATION AND IMPROVEMENT (Study 2)</b></p> <ul style="list-style-type: none"> <li>- Assumptions about the system's role</li> <li>- Assumptions about the HAI issue</li> </ul>	<p><b>OP1: ORGANISATIONAL ASSUMPTIONS DRIVING AUTOMATION IMPLEMENTATION AND IMPROVEMENT (Study 3)</b></p> <ul style="list-style-type: none"> <li>- Assumptions about the system's role</li> <li>- Assumptions about the HAI issue</li> </ul>
--	--

### 7.2.2. SME feedback on OP2: Organisational capability for handling HAI

This theoretical category suggests that the handling of HAI issues requires the presence of a focused organisational capability—this capability including a specific process, the valuing of end-users ‘involvement, i.e., the controllers, the availability of supporting software tools and organisational roles (§ 6.4.2).

#### Positive attitude towards air traffic controllers’ involvement

The experts seemed supportive of the idea that valuing the involvement of the end-user—the air traffic controller—is an important precursor to the correct implementation of safety nets. To support this point, they noted that best in class implementers do not limit end-user involvement to the pre-implementation phase of the safety net lifecycle, but extend it also to operation—the phase in which safety nets become available in the control room. In other words, the consideration given to the end-user perspective is reflected also in the way service providers establish mechanisms for collecting systematically air traffic controllers’ feedback related to operational safety nets after they enter operation.

This should not be taken for granted, as some providers essentially may lack effective reporting arrangements. Experts, in particular, stated that organisational contexts can inhibit controllers’ reporting of known safety nets related problems. This may be due to one or a combination of the following conditions:

1. *Reports not triggering corrective actions.* One expert (SME2) reported a service provider in which controllers were asked to provide a written report of each unwanted, missing or late safety net alert. However, such reports remained unaddressed by management, and after an initial period, the controllers eventually stopped reporting, as they began to view this activity as an additional and unnecessary activity. On the other hand, management interpreted the lack of reporting as a sign that the safety nets were successfully tuned, and that no improvement action was needed.
2. *Hierarchical command and control culture.* Also, controllers’ reporting could also be inhibited by the hierarchical culture that might be found in some ANSPs. Such a culture may inhibit open discussions about safety nets issues without approval from the relevant leaders. This is especially true for people external to the organisation. This observation was made by a consultant with experience in safety net implementation in several countries.
3. *Blaming controllers for safety net problems.* Finally, it was mentioned that senior management might interpret nuisance alerts as an indication of poor controller’s

performance, rather than as a sign of problems with safety nets (SME1; SME2; SME7).  
For instance:

*“We were asked to study the behavior of STCA in a centre. Indeed, some nuisance alerts were noticed. Not a real point, we thought, but we gave advices on how to improve the STCA tool. Afterwards we understood that the real point was that management thought that every STCA alarm was a mistake of the controller.” (SME1).*

Overall, these conditions may favour patterns of organisational inertia regarding safety nets: the hazards and problems with these systems may remain unaddressed in the organisation despite their (negative) impact on the activity of the sharp end practitioner. When compared to the early definition of the framework, these considerations are relevant because, on the one hand, they stress the confidence of the importance of the service provider valuing the end-user perspective; and on the other they suggest that this is important not just during the implementation phase, but also during the operational life of these systems. Valuing the controllers’ involvement seems important to ensure that emerging problems are identified and then addressed.

#### Senior management support to the safety nets domain

SMEs’ reactions to the initial definition of the framework suggested that while this captured adequately the importance of the availability of a parameterisation process, it seemed to emphasise this item too much compared to two important organisational conditions: the availability of senior management support, and the availability of a dedicated safety net team. In the SME’s views these were two important elements of what makes an organisational capability, which however were not adequately represented by the OP2 sub-category “Availability of supporting software tools and a structure roles”. The importance of senior management support, and the availability of a supporting dedicated team are addressed in the reminder of this section and the next section respectively.

The experts considered it very important to have senior management committed to safety nets implementation and continuous improvement. Senior management commitment was defined as:

*“...allocating the appropriate resources to the improvement of the overall operations of safety nets, [so that] safety nets are directly allocated under certain budget and manpower” (SME4).*

In particular, senior management commitment was considered important to:

1. Ensure a long term sustained organisational focus over the domain of safety nets and the accompanying problems. Implicit in the quote above is the conviction that continual senior management support is essential to ensure continuity in the quality of safety nets implementations.
2. Promoting the development and retention of expertise. It takes time and effort to establish a more controlled and repeatable process: safety net introduction is not, in fact, a “one time exercise”; it is, instead, a continuous learning process in which the organisation learns over time to better control the problems associated with these systems (SME1, SME3, SME4, SME6). One expert reported that for his organisation it took considerable efforts to move from an initial, problematic, trial and error parameterisation to a more controlled one, i.e., one that delivered safety nets that were considered acceptable by controllers (SME9). To make this kind of incremental learning possible it is essential that senior management “nurture the expertise” in the safety nets domain in the long term. Commenting on expertise, it was noted that senior management long term support is important especially for ANSPs with a high staff mobility to ensure the retention of safety net expertise. Safety net expertise “needs to be retained also across different generations of developers, as an ANSP might encounter several safety net implementations [over time], so you do not want to lose the lesson learnt from the previous implementations” (SME4).
3. Promoting awareness of safety nets’ role within the organisation. Senior management support was also considered essential to raise awareness of the purpose of safety nets across the different hierarchical levels and functional units of the organisation (SME4). As stated by an expert:

*“Senior management should spread awareness throughout the organisation, both vertically and horizontally, insist on the unique understanding of the purpose and the role of the safety nets within the company's ATM system and commitment of all the segments of the company to supporting this role.” (SME4).*

To illustrate this point, the case of one organisation was cited, in which the staff involved in the maintenance of a piece of equipment would regularly render a safety net unavailable. Not being directly involved in the provision of Air Traffic Services, the technical training part of the organisation was not aware of the role of the safety net in the overall system, and was pursuing their own internal needs and priorities. Eventually, it took the intervention of senior management to clarify the importance of safety nets across the different divisions of the organisation. In particular, it mandated the

development of a dedicated internal safety net policy. This was intended to make explicit what safety nets are and how they should be managed. In this particular instance, such a policy was instrumental in clarifying that safety nets performances should take priority over other operating constraints, and should be compromised only if justified by safety requirements (SME3).

#### Safety nets team

It was also noted that it is essential that senior management creates a permanent multidisciplinary team within the organisation, dedicated to the implementation, and continuous monitoring and improvement of safety nets. Such a team must include as a minimum one safety net engineer, an air traffic controllers' representative, and should be led by a professional project manager (SME3; SME5; SME10). The presence of such a team is essential in order to ensure that the ANSP's resources are appropriately coordinated in order to solve safety net problems. This was the lesson that the experts had learnt after they had experienced directly problematic implementations having poor internal coordination as a precursors, e.g.:

*"To make my point short: in this service provider, operational, technical and safety expertise did not work together to solve problems. It turns out that after our (EUROCONTROL) intervention and analysis, the organisation has learned that working together can solve problems... Introduction should be driven by a multidisciplinary team (operational, technical, safety expertise)." (SME1)*

*"Fortunately we have learned from the past and new safety nets or controller assistance tools which were developed together with users (ATCOs), product management, requirements and software engineers... [Working] together it is always much easier, very often more successful and at least accepted by the users!" (SME3)*

It was noted that the absence of a multidisciplinary team dedicated to safety nets may be observed especially in small service providers: these may lack the resources and manpower to have staff permanently allocated to the safety net domain. One expert, for instance, witnessed a small ANSP in which one software engineer alone, not a team, was held responsible for the implementation of multiple safety nets (SME5). Similarly, another expert described the experience of a small provider in which an operationally active controller was also assigned, on a part-time basis, the project management of safety net implementations (SME10). Eventually this person was unable to follow the implementation. Neither situation points at limitations in the competence levels of the individual concerned, or of personnel of small service providers in general. Instead, they

emphasise the fact that managing the intricacies of safety nets implementation requires the integration of different areas of expertise.

#### Safety nets specific performance framework

Another item missing from the initial definition of the emerging framework of organisational precursors was the presence of a safety net specific performance measurement framework, i.e., a set of metrics or indicators specifically defined for monitoring the performances of the MSAW and other safety nets. Such a framework permits personnel to monitor areas of safety nets performance (also based on statistical data), and decide which area is in most in need of improvement (SME6; SME4). While this framework is essential to ensure continuous monitoring and improvement of safety nets. Lack of such a performance framework limits the ability of service providers to monitor how well they are managing nuisance alerts, as well as other safety nets related problems, e.g.,:

*“Several years ago, \*\*\*[name of the service provider omitted] did an assessment of their safety nets. They noted that they had 3000 alerts in a week—and concluded that the performance was ‘OK’. This is a self-assessment with no reference or comparison to anyone else’s system.” (SME5)*

The performance framework needs to be specifically designed for safety nets to be effective. In support of this point, one expert noted that currently service providers have to comply with a global target safety level (TSL), which is expressed as the tolerable number of separation loss per year. However, while this indicator provides an aggregated measure of the overall level of safety achieved by the organisation, it only provides an indirect measure of safety net performances, as there are in fact several additional areas of operation that might contribute to the TSL, rather than safety nets alone. As a result, while focusing on this global indicator, senior managers might fail to see all the intricacies involved in ensuring safety nets operational fit (SME4).

#### 7.2.2.1. Implications for OP2

While connecting to the experiences of the participating SMEs and triggering supportive feedback, the presentation of the second precursor of the emerging framework, OP2: “Organisational capability for handling HAI issues”, also triggered some constructive criticism. This was directed towards the first sub-category of OP2, “Internal parameterisation”, suggesting this sub-category was reasonably represented in the initial definition of the framework—at least represented in a way that does not trigger negative



reactions. It was the second and the third sub-categories of OP2—“Valuing of air traffic controllers involvement” and “The availability of supporting software tools and safety nets specific structure of roles”—that attracted the majority of the SMEs’ comments. Regarding the former, it was suggested that the proper consideration of the end-user is not limited to the implementation phase, but extends to the proper reporting of safety nets issues. Having an environment that does not promote the reporting of safety nets promotes in fact patterns of inertia over the improvement of these systems. Regarding the second item, experts’ feedback emphasised the role of senior management support and the availability of an internal multidisciplinary team permanently dedicated to safety net implementations. These elements were deemed to be important components of a safety net capability that should stand alone. Finally, still commenting on the third item of the framework, it was noted that another important element is the availability of a safety net specific performance framework—a set of metrics specifically defined to capture issues with the safety nets. To summarize, one item of OP2 was confirmed as is, some were refined and expanded, and another element was suggested as an addition. Overall, this feedback allows for the definition of a more expanded and sounder version of OP2, as reported below in Table 22.

*Table 22. Initial (left box) and updated version (right box) of OP2 as resulting from Study 3.*

<b>OP2: ORGANISATIONAL CAPABILITIES FOR HANDLING HAI ISSUES (<i>Study 2</i>)</b>	<b>OP2: ORGANISATIONAL CAPABILITIES FOR HANDLING HAI ISSUES (<i>Study 3</i>)</b>
<ul style="list-style-type: none"> <li>-Parameterisation process</li> <li>-Positive attitude towards air traffic controllers involvement</li> <li>-Availability of supporting software tools and a structure of roles</li> </ul>	<ul style="list-style-type: none"> <li>-Parameterisation process</li> <li>-Positive attitude towards air traffic controllers’ involvement</li> <li>-Senior management support</li> <li>-Supporting tools</li> <li>-Safety nets team</li> <li>-Safety nets specific performance framework</li> </ul>

### 7.2.3. SME feedback on OP3: Control over implementation quality at the boundary between the service provider and the software manufacturer

This precursor captures the fact that the interaction between the ANSP and the software manufacturer supplying and installing the MSAW can have a bearing on the implementation of the system: on the one hand the service provider may have the knowledge to specify the requirements of the system it is purchasing, on the other hand, the organisation may inadvertently “transfer” control over the quality of the purchased system to the manufacturer, thus risking being provided with a suboptimal system (§ 6.4.3). The following categories capture the reactions of the participating SMEs to OP3:

1. **Pattern common to other less-successful implementation contexts.** The experts confirmed that they had experienced different situations in Europe in which the service provider purchased the MSAW and other safety nets without specifying the requirements of these systems. Similarly to the situation described for the Deltasky case, the service provider was then provided with safety nets which were not fit for purpose, as they were not specifically parameterized for the specific operational environment, e.g.,:

*“Some ANSPs have enough knowledge to work closely with the manufacturer...However, in very many cases, the manufacturer delivers a system “as is”, with no input from the ANSP. Many ANSPs leave it to the system supplier to set up the parameters.” (SME5).*

This expert reported to having observed this dynamic in at least three ANSPs and in one military control centre. There, the implemented MSAW and STCA turned out to be poorly suited to the specific operational context of the client organisations. Specifically to the MSAW, this expert noted that the manufacturer spent insufficient time on setting up an MSAW terrain data base (SME5). The same incident was witnessed by two other experts (SME2, SME7), who added that the organisations in question did not consider the issue of parameterisation before buying their safety nets. One expert (SME7) confirmed that this *“is not [the best strategy] to obtain the best service from the manufacturer. There is the risk that the manufacturer will try to obtain the maximum result with the minimum effort.”* (SME7). All of the three organisations, eventually, appreciated the significance of the problem of nuisance alerts only when the system went operational. This suggests that the dynamic observed in Deltasky was not limited to that service provider, others may be affected.

2. **Problem affecting also large ANSPs.** In particular, although the Deltasky case might

suggest that this problem is limited to small service providers, it was noted that this is not the case. The risk has been reported to occur also in the case of large service providers that regularly purchase their systems—not just safety nets—from their national software manufacturer (SME7, SME8). In these cases, it was observed that the software manufacturer operates in a near monopolistic market, which may favour more the commercial interest of software manufacturer versus the operational needs of ANSP, e.g., *“I think monopoly is the main problem. With monopoly, [\*\*\*name of the ANSP omitted] operates any system that [\*\*\*name of the software manufacturer omitted] will pass. Nothing can change at the level of the Human Machine Interface unless [the ANSP] can purchase its systems from other software manufacturers”* (SME8). This comment was made by an expert commenting on the poor safety net HMI of a large South European service provider which regularly purchases its system from its national software manufacturer.

3. **Working closely with the manufacturer to avoid the concerned problem.** When reflecting on how to avoid situations such as the one experienced by Deltasky, two considerations have emerged. First, the obvious solution to the problem would be to specify formal safety nets requirements at the beginning of a project, so that these can be included in the purchasing contract—with this contract clearly defining the responsibilities that the manufacturer will have over the system during the operational life of the system (SME6).

At the same time it was acknowledged that—and this leads to the second consideration—it might be very difficult for a service provider to define upfront safety nets requirements when facing its first implementation—that is, when no previous experience with the development and tuning of the safety nets is available within the organisation, e.g.:

*“To say the truth, it’s difficult for an ANSP to develop a clear and elaborate concept of operation in a vacuum of experience.”* (SME4).

Essentially, at the time of their first implementation, safety nets, like any other system, are seen as an innovation, i.e., as something that requires a period of familiarisation and trial and error before the organisation possesses the knowledge to specify an elaborate a concept of operation (SME8). In the absence of such knowledge, it was suggested that the purchase and implementation of safety nets should be framed in the context of a *“continuous partnership [between the manufacturer and the ANSP] to allow for safety nets operations to evolve in [the] specific operational environment.”* (SME4.) Such a continuous partnership could be initiated by ensuring appropriate expertise overlap between the ANSP and the safety nets manufacturer, e.g.:

*“If you have to buy a novel safety net, or you want to upgrade your*

*system, it is better to choose the supplier that tells you: “I will spend a period in your control center and will teach your experts how to parameterize the system.” (SME7)*

This expert also stressed the importance of establishing an initial collaboration phase between the ANSP and the manufacturer during which requirements are jointly explored and defined. Essentially, the point is that during the first implementation of safety nets, the so-called “black box approach”—the classic outsourcing approach in which system requirements are specified right at the outset—does not work, because both the service provider and the manufacturer need to have mutual access to their expertise in order to develop the system. The ANSP brings knowledge of its unique operational context, and this knowledge is necessary for the fine tuning of the safety nets; the software manufacturer provides software knowledge, which is necessary to build the system.

### 7.2.3.1. Implications for OP3

The feedback collected for the third category of the emerging framework—*OP3 Control over implementation quality at the service provider-software manufacturer boundary*—essentially confirmed the role of this precursor in explaining the quality of human machine interaction as found in the operational automated system. In particular, three corroboratory categories have emerged. Categories 1 and 2 indicated that the problems involved in the trust based control observed in the Deltasky case were not limited to this organisation, but were also observed in other service providers—both small and large ones. Finally, the last category of feedback, in addition to further confirming the accuracy OP3, also expanded this category, for it highlighted the fact that the actual control over implementation depends on the prior knowledge of the service provider. Without such a knowledge, it is important to frame the relationship with the software manufacturer as a long term partnership. Overall, this feedback seems to confirm the influence of OP3 in the handling of HAI issues within safety critical service provider organisations.

**Table 23.** Evolution of OP1 as a result of Study 3: this category has remained unchanged.

<p><b><u>OP3: CONTROL OVER IMPLEMENTATION QUALITY AT THE SERVICE PROVIDER MANUFACTURER BOUNDARY (Study 2)</u></b></p>	<p><b><u>OP3: CONTROL OVER IMPLEMENTATION QUALITY AT THE SERVICE PROVIDER MANUFACTURER BOUNDARY (Study 3)</u></b></p>
---	---

### 7.3. SUMMARY AND DISCUSSION

---

The aim of this study was to explore the plausibility of the organisational patterns of the initial version of the OPHAI framework resulting from Studies 1 and 2. By means of a focus group discussion combined with a qualitative questionnaire, this study obtained corroboratory feedback from a group of 11 safety net experts. These experts were invited to comment on the definitions of the identified organisational precursors of the framework, based on knowledge coming from their direct and indirect professional experience.

The evidence collected in this study took the form of a total of 11 corroboratory categories: four categories identified for OP1; four, for OP2; and three, for OP3. It can be noted that some differences can be observed regarding the relative depth of these categories. The corroboratory categories identified for OP1 and OP3 are organised on one hierarchical level; those identified for OP2, on the other hand, are organized on two hierarchical levels. In other words, the corroboratory categories identified for OP2 enjoy a deeper structure than OP1 and OP3. This difference reflects the fact the experts had different types of reactions to OP1 and OP3 on the one hand, and OP2 on the other. OP1 and OP3 received essentially supporting feedback, where the experts' response to the definitions of these categories of organisational precursors mainly suggests that they were witnessed by the experts participating in the study (as concluded under § 7.2.1.1 and § 7.2.3.1). So, this expert's response added confidence without proposing a refinement. Somehow, this suggests that the earlier definition of OP1 and OP3 was perceived as trustworthy.

The earlier definition of OP2, on the contrary, was perceived as being too narrow in scope. The definition of organisational capability for handling HAI issues, in fact, was perceived as underrepresenting important components, namely the role of senior management, the existence of a safety net team, the existence of a dedicated safety net performance framework (as concluded under § 7.2.2.1 ). According to the experts, these elements needed to be more visible in the framework, and therefore an update and more trustworthy version of OP2 had to include them.

Having elaborated on how the collected evidence evolved the OPHAI framework, it becomes important to elaborate on the implications of the same evidence on:

- *OPHAI's interpretive validity*<sup>7</sup>, i.e., the extent to which the identified categories of organisational precursors capture organisational dynamics that appeared plausible according to the viewpoints and experiences of the interviewed experts;

---

<sup>7</sup> See also § 4.3.4 for a definition of interpretive validity.

- *OPHAI's generalisability*, the extent to which the framework can be applied to other organisational contexts besides those of Alphasky and Deltasky.

Regarding interpretive validity, it can be said that the study has allowed to collect corroboratory evidence which appeared to improve the initial OPHAI's version (as delivered by Studies 1 and 2). In particular, the study has identified a total of 11 corroboratory categories, which capture relevant experts' opinions, experiences, and suggestions for improvement about the framework. This evidence suggests that the OPHAI's categories of organisational precursors have been witnessed by the participating experts in real life situations, i.e., situations related to the implementation and improvement of the MSAW and other safety nets occurred in European ANSPs. The fact that these patterns resonate with the knowledge of the selected experts is an important indicator of the study's interpretive validity. Two reasons support this position. In qualitative research member checking is usually accepted as one of the most important validity check (Morgan, 1997) (§ 4.3.4). Furthermore, the importance of member checking is especially important in the case of social and organisational phenomena that, as for the phenomenon under study in this research, cannot be directly observed, or measured by objective and independent quantitative data collection means. In this case recurring to the opinions of experts become a necessary means to draw conclusions.

Regarding generalisability of the framework—i.e., the extent to which the results can be applied, at least, to other service providers—it can be noted that the expert professional experience was also based on contexts other than Alphasky and Deltasky. This invites some scepticism about the idea that the identified categories of precursors are the products of the unique organisational contexts of Alphasky and Deltasky. Rather, it seems reasonable to expect that the same precursors may occur also in other ANSPs, at least European ones. In addition to the collected experts' view, one further consideration can support this point. In particular, European ANSPs operate in a common competitive and institutional environment, i.e., the European air traffic management system, which dictates common European-wide modernization roadmaps, and common rules, standards and policies for the conduct of commercial practices and safety. Arguably, these conditions expose European ANSPs to similar common pressures for modernization, cost-reduction, and productivity.

Of course these considerations do not address the question of whether the framework can be generalised to other contexts other than European ANSPs. This important question will be addressed in the next chapter, together with a reflection on the broader theoretical implications of the framework for the understanding of HAI issues in in safety-critical domains.

## 7.4. CHAPTER CONCLUSIONS

---

This study reported on the third and final empirical study of the thesis. The study aimed to corroborate and refine the initial version of OPHAI delivered by Studies 1 and 2. The study consisted of an SME study based on a sample of 11 safety net experts. The evidence collected further confirmed the first and the third organisational precursors of the framework, while the second precursor had to be refined in order to better reflect the experiences of the experts. This evidence led to a refined version of the framework, which is reported below. At this point, what remains to be discussed at this point are the broader implications of the framework for the understanding of HAI issues in complex, high consequence organisations, and what are the practical implications of the framework. These questions will be addressed in the following chapter.

*Table 24. The revised and expanded version of OPHAI, as resulting from Study 3. While OP1 and OP3 have been confirmed, OP2's subcategories have been refined and expanded.*

<p><b><u>OP1: ORGANISATIONAL ASSUMPTIONS DRIVING AUTOMATION IMPLEMENTATION AND IMPROVEMENT</u></b></p> <ul style="list-style-type: none"> <li>- Assumptions about the system's role</li> <li>- Assumptions about the HAI issue</li> </ul> <p><b><u>OP2: ORGANISATIONAL CAPABILITY FOR HANDLING HAI ISSUES</u></b></p> <ul style="list-style-type: none"> <li>- Parameterisation process</li> <li>- Positive attitude towards air traffic controllers' involvement</li> <li>- Senior management support</li> <li>- Supporting tools</li> <li>- Safety net team</li> <li>- Safety net specific performance framework</li> </ul> <p><b><u>OP3: CONTROL OVER IMPLEMENTATION QUALITY AT THE BOUNDARY BETWEEN THE SERVICE PROVIDER AND THE SOFTWARE MANUFACTURER</u></b></p>
--

# Chapter 8.

## **DISCUSSION and CONCLUSIONS**

### **8.1. CHAPTER INTRODUCTION**

---

To recap: the main aim of this thesis was to inquire into the organisational precursors to human automation interaction (HAI) issues that can be found in safety-critical domains. The theoretical motivation behind this inquiry was to explore the theoretical landscape lying at the boundaries of the current main perspectives on the problem, namely the human computer interaction (HCI) and system lifecycle ones (chapter 2), and the enlarged organisational safety (OS) perspective (chapter 3). As a result of this effort, the thesis has delivered the organisational precursors to human automation interaction issues (OPHAI) framework. This framework is empirically grounded, because it was developed in a bottom-up fashion based on evidences collected from three qualitative empirical studies: two retrospective, qualitative case studies, and a third subject matter expert study. The first study (chapter 5) identified the framework's first component, OP1, based on an analysis of the NTSB-issued safety recommendation letters and the safety recommendations that addressed the MSAW, and the relevant correspondence exchanged between the NTSB and the FAA. Informed by the organisational experiences of both a successful and a less successful MSAW implementer, the second study (chapter



6) identified the second and the third main categories of organisational precursors to HAI issues, OP2 and OP3; and provided additional support to OP1. The third study (chapter 7) expanded and refined the OPHAI framework. It collected additional corroboratory evidences in support of OP1 and OP3, and delivered a more refined version of OP2. Table 25 summarises the contribution of the three empirical studies to OPHAI's development. Table 26 in the next page provides the generalised version of the framework.

**Table 25.** Contribution made by the empirical studies to the OPHAI framework's development.

<b>OPHAI's components</b>	<b>Study 1</b> (chapter 5)	<b>Study 2</b> (chapter 6)	<b>Study 3</b> (chapter 7)
<b>OP1</b>	Identified OP1 (§5.4)	Collected supporting evidence for OP1 (§6.4.1)	Collected supporting evidence for OP1 (§7.2.1)
<b>OP2</b>		Identified OP2 (§6.4.2)	Collected evidence for refining OP2 (§7.2.2)
<b>OP3</b>		Identified OP3 (§6.4.3)	Collected supporting evidence for OP3 (§7.2.3)

The remainder of this chapter is organised as follows:

- Sections 8.2 and 8.3 compare the framework with the relevant literature reviewed in chapters 2 and 3;
- Section 8.4 describes the potential practical uses of the OPHAI framework. In order to highlight the framework's distinctive features, this section also includes a comparison with relevant comparable models, i.e., models that either address HAI issues or consider organisational precursors to failure;
- Section 8.5 addresses the generalisability of OPHAI;
- Section 8.6 addresses the limitations of the research and opportunities for future research;
- Section 8.7 concludes by summarising the contribution to knowledge the thesis has made.

*Table 26. Generalised version of the OPHAI framework.*

<p><b><u>OP1: ORGANISATIONAL ASSUMPTIONS DRIVING AUTOMATION IMPLEMENTATION AND IMPROVEMENT.</u></b> The assumptions held by management and developers during management and implementation. Two types of assumptions can be distinguished:</p> <ul style="list-style-type: none"> <li>- <b>Assumptions about the system's role</b>, i.e., how and for which purpose the system should be used by the end-user, and for which purpose, and how it will benefit the organisation;</li> <li>- <b>Assumptions about the HAI issue(s)</b>. Assumptions of developers and management about the nature and severity of the HAI issues that may emerge during the use of the automation.</li> </ul> <p><b><u>OP2: ORGANISATIONAL CAPABILITY FOR HANDLING HAI ISSUES.</u></b> The set of (interrelated) organisational routines that allow the organisation to reliably handle HAI issues. The components of such an organisational capability include:</p> <ul style="list-style-type: none"> <li>- <b>Parameterisation process</b>. The process needed to parameterise the automation in order to adapt it to the specific context in which it will be used. This process is repeated whenever the organisation implements the automation at a new site;</li> <li>- <b>Attitude towards end-user involvement</b>. The attitude towards the continuous systematic involvement of end-users in the implementation and improvement of the automation, and the positive consideration of the HAI issues end-users may report;</li> <li>- <b>Senior management support</b>. The continuous support of senior management to the automation domain, as manifested in the automation policies, and the resources allocated to automation management and its continuous improvement;</li> <li>- <b>Automation team</b>. The team of staff members permanently allocated to the management of automation and its continuous improvement;</li> <li>- <b>Supporting tools</b>. The set of software tools (e.g. demonstrators, recording tools, etc.) developed in the organisation to monitor and improve the quality of the automation;</li> <li>- <b>Automation specific performance framework</b>. The framework of HAI issues specific performance indicators.</li> </ul> <p><b><u>OP3: CONTROL OVER IMPLEMENTATION QUALITY AT THE BOUNDARY BETWEEN THE SERVICE PROVIDER AND THE SOFTWARE MANUFACTURER.</u></b> The type of control mechanism(s) established to ensure that the manufacturer will act in the service provider's best interest—i.e. ensuring that the automated system purchased is actually suitable for the operational context of the service provider.</p>
--

## 8.2. THEORETICAL CONSIDERATIONS

---

### 8.2.1. Comparisons with the HCI perspective

An important area of study that has dealt with the problem of HAI issues is the Human Computer Interaction (HCI) perspective (§ 2.2). Notably, this perspective unites scholars that have elected as their object of study the unit composed of (i) the human (i.e., the user(s) of the technology), (ii) the automated system, and (iii) the unfolding interaction between (i) and (ii). The conceptualisation of HAI issues is viewed as arising from the analysis of these elements (although from very different research traditions, such as human information processing (§ 2.2.1), distributed cognition (§ 2.2.2), activity theory (§ 2.2.3), computer supported collaborative work (§ 2.2.4), and cognitive system engineering (§2.2.5)). Notably, HCI models provide a set of concepts extremely useful for designers, human factors and safety practitioners and researchers to diagnose HAI issues related to a given automated system. In turn, this provides the foundations for further formal and informal evaluations, and corrective design changes (§ 2.2.6). In comparison, OPHAI points out a set of organisational precursors to HAI issues. Such precursors are exogenous to the unit of analysis as defined by the HCI models. In fact, these precursors are not visible from the perspective of the user of the technology, as they are located at organisational levels higher than the level of operations. Yet, they are important to consider if one, in addition to identify HAI issues, also wants to know how a given service provider may (inadvertently) have left such issues unaddressed in the operational system or it has succeed in managing them. In other words, OPHAI's precursors are important to consider if one wants to better understand the organisational realities involved in the management of safety-critical automation.

In addition to this, another consideration arises from the comparison of OPHAI and the HCI perspective. In fact, a basic idea found in this perspective is that the introduction of novel automation introduces new task demands for the front-end operators of complex systems (e.g., Woods, Dekker, Cook, Johannesen, & Sarter, 2010; Sarter, Woods, & Billings, 1997; Woods & Sarter, 2000). OPHAI resonates with this idea, in that its dimensions suggest that novel automation brings new demands not only for the automation's end-user, as usually discussed in the HCI perspective, but also for the organisation that has to implement and manage the automation. In fact, OPHAI suggests that the organisation needs to direct resources to the development of adequate assumptions about the system's operational role and the nature and significance of the accompanying HAI issues. This is because such assumptions appear to influence the kind

of organisational responses directed towards the handling of HAI issues (§ 5.4.2, § 6.4.1, and § 7.2.1). In addition, the organisation needs to establish adequate organisational capabilities for handling HAI issues, capabilities that include the availability of a parameterisation process; adequate involvement of the end-user; continuous support from senior management; the availability of a permanent team dedicated to the management and improvement of automation; and the availability of an automation-specific performance framework (§ 6.4.2 and § 7.2.2). Finally, the organisation needs to establish adequate control mechanisms in order to ensure control over the quality of the system supplied by the software manufacturer (§ 6.4.3 & 7.2.3). Overall, these organisational conditions warn policy makers, regulators, managers and developers of safety critical service providers that the introduction and operation of novel automation is not a sole technical or engineering effort. Automation comes with organisational demands that need to be addressed in order to control the HAI issues accompanying it.

Furthermore, these demands are not limited to the implementation phase of the automation, but extend across its operational life. This point is best supported by OP2, because the availability of a dedicated organisational capability for handling HAI issues (found in the successful case in Study 2, and later refined and expanded in the Study 3) reflects organisational conditions that are permanently established in the organisation in order to ensure the continuous monitoring and improvement of the automation. These considerations are relevant because, although stakeholders at the blunt end usually acknowledge the very intense effort needed to put a new technology into use (e.g., Humphreys et al., 2006), they do not necessarily see the resource demands that are needed past the “O-date” to sustain the new technology (Campbell, Sittig, Ash, Guappone, & Dykstra, 2006).

### **8.2.2. Comparison with system lifecycle perspectives**

The value of the OPHAI framework can also be understood by comparing it with the system lifecycle models (§ 2.3)—such as user centred design (UCD) (ISO, 2010), human factors integration (HFI) (HFI DTC, 2007), and system safety (e.g. Roland & Moriarty, 1990). The framework highlights, in fact, some organisational dimensions that are not considered by these models. For instance, UCD prescribes user requirements identification, by means of methods such as ethnographic studies and usability evaluations in order to achieve increased system usability. HFI promotes various human factors analysis along the system development lifecycle, in order to make sure that human factors aspects are systematically captured and addressed during the engineering lifecycle

of a system. System safety pursues similar objectives in order to ensure that safety hazards and requirements are appropriately identified and addressed during the engineering lifecycle of a system. In maintaining this focus, these approaches provide useful process models which can explain the quality of human automation interaction found in safety-critical organisations mainly in terms of adherence to or deviations from a predefined set of process guidance. However, these models do not consider the organisational factors at the blunt end, highlighted by OPHAI, which may affect the way HAI issues are addressed. These models neglect these factors because their focus is mostly on the process aspects.

### **8.2.3. Comparison with the OS perspective**

The primary distinctive trait of the OPHAI framework presented here is that of highlighting some of the organisational dimensions involved in the handling of technology's side effects, i.e., HAI issues. Essentially, the framework's categories support the idea that HAI issues can be seen as the symptom of deeper organisational phenomena, i.e., they can be traced back to specific organisational dynamics and conditions at the blunt end of the organisation. In general, this orientation receives support from models in the OS area (reviewed in chapter 2), as these models also trace the potential for success and failure in organisational (safety) performances to dynamics and conditions found at the blunt end of the organisation.

In particular, existing OS models primarily support the first category of organisational precursor (OP1), which asserts that the handling of HAI issues in safety-critical domains stems from the organisational assumptions driving the implementation and improvement of automation. Scholars from the OS area (as well as from classic management) have long stressed the role of shared organisational interpretive frames in driving organisational responses: it is the interpretation of the world, i.e., the shared cognitive frames of organisational actors that drive organisational responses to signs of danger and ambiguous anomalies (Weick & Sutcliffe, 2011; Vaughan, 2009; Edmondson et al., 2005; Milliken, Theresa, & Bridewell-Mitchell, 2005; Turner & Pidgeon, 1997; Demchak, 1991). When this interpretation of the world and risk is inconsistent with the way the world really is, the potential for organisational failure increases as the organisation accepts more and more risk than it can actually manage. While these considerations highlight the link between organisational frames at the blunt end and safety performances, this research suggests that a similar link exists with regard to the handling of the undesired effects of automation technology.

However, OPHAI highlights two theoretical constructs, OP2 and OP3, that are not usually discussed in models from OS area, which appear however to be relevant to understanding the handling of HAI issues in safety-critical industries. In particular, the idea that handling HAI issues requires the availability of a dedicated organisational capability (OP2) coheres with models in the innovation and business performance literature (Leinwand & Mainardi, 2010; Leonard-Barton, 1998; Prahalad & Hamel, 1990). These models emphasise that specific organisational performances reflect the unique combinations of tools, skills, processes, values, supporting roles, and leadership that are available in organisations. It is the availability and combinations of these factors—i.e., the element of an organisational capability—that allows an organisation to achieve and sustain certain organisational outcomes over time. In this research, the notion of organisational capability provided a comprehensive construct with which to characterise the distinctive elements, or practices, activated in the organisation to handle the HAI issue in question.

Analogous considerations apply to OP3. The idea that the problematic handling of HAI issues can be traced back to the relationship between a safety-critical service provider organisation and the software manufacturer receives support from the IT outsourcing literature (Tiwana, 2004; Willcocks & Sauer, 2000). This literature acknowledges that one important challenge organisations face in outsourcing software is to ensure that the software manufacturer delivers as expected (Choudhury & Sabherwal, 2003; Kirsch et al., 2002; Kern & Willcocks, 2001). While this view was developed based on the analysis of non-safety critical applications, this research also stresses the importance of control at the organisational boundary between software manufacturers and safety critical service providers.

### **8.3. METHODOLOGICAL CONSIDERATIONS**

---

In addition to the theoretical considerations outlined above, two important methodological considerations can be made. One highlights the originality of the data used in Study 1; the second, the value of the notion of “anomaly trajectory in the organisation”, which was used in Studies 1 and 2. Note that these considerations arise from a comparison with the models developed in the OS area. This is because the research design adopted in this study is essentially derived from the single case study approach, which is frequently used in the OS area (§ 4.3).

#### **8.3.1. Considerations on the data used in Study 1**

Study 1 (chapter 5) made use of original data sources, namely the safety recommendations of the US official accident investigation agency (the NTSB), the letters

justifying and conveying these recommendations, and the response letters by the recipient organisation (the FAA) (see § 4.4.1.1 for the description of these sources). Prior to this research, only Tasca (1990) was found to use this type of data, although for a different purpose: inquiring into the broader organisational precursors to human error in the maritime domain. This study's insights suggest that the same type of data could be used for a comparable research objective, namely investigating the organisational precursors to HAI issues. This did prove to be the case as the type of data in question made possible to identify the first category of organisational precursor of the OPHAI framework. This consideration is important because further studies can exploit the same type of data to investigate the organisational trajectory of types of HAI issues other than the one investigated in this study. Furthermore, the same data can be used to investigate the organisational trajectory of other safety issues—not necessarily automation related—to explore the ways in which these have been (mis) handled by the organisation.

In suggesting the possibility of using safety recommendations, safety recommendation letters, and ensuing correspondence letters it is important to elaborate on the relative advantages and disadvantages of these data sources. Their primary advantage is their availability in the public domain. In other words, this data do not require the negotiation of organisational access. These sources can be retrieved from the relevant on-line database(s). This advantage is not negligible, considering that in OS research the negotiation of organisational access is usually considered a “major hindrance” (Bourrier, 2011). Thus, this kind of data can be an addition to the range of publicly available data sources commonly used in the OS area—i.e., official accident and incident investigation reports.

However, two disadvantages of using these sources can be mentioned. Firstly, using original data may entail a relatively long familiarisation stage. In Study 1, nearly two years elapsed before the researcher was able to identify a plausible data analysis strategy. (As described in chapter 3, it was only after the familiarisation stage that it became clear that the body of data included three sub-sets of data.) It must be noted that this difficulty results not from the novelty of the type of data considered, but also from the exploratory nature of the research.

The second disadvantage concerns data representativeness. Safety recommendations, safety recommendation letters, and the ensuing written exchange represent essentially the official view of the organisation with regard to specific safety problems and HAI issues. Therefore, this should be coupled with interviews of organisation insiders to deepen the

rationales and perspectives that have led to the development of such a view. In this research, this limitation of Study 1 was addressed by including in the two subsequent studies interviews with relevant personnel and experts. Prospective researchers interested in using similar data for organisational research will have to consider this limitation, and the need to mitigate it by either complementation or triangulation with other data sources.

### **8.3.2. Considerations on the value of the organisational trajectory of an HAI issue**

The second most distinctive trait of the case studies used in this research (Studies 1 and 2) was their focus on the organisational trajectory of an HAI issue related to a specific technology, the MSAW. The case studies focused on the history, or lifecycle, of one type of anomaly, an HAI issue, faced by the organisation. This analytical focus was inspired by Vaughan's seminal investigation of the Challenger Space Shuttle disaster (Vaughan, 2009, 2004, 1997). In that work, Vaughan traced the historical trajectory of an anomaly within NASA—namely, the O-ring erosion problem—with the intention of understanding how deviations from accepted and safe norms were gradually accepted by NASA's management until the disaster unfolded.

Regarding the value of an analytical focus on the anomaly trajectory in organisations, this research provides two key insights in addition to those provided by Vaughan's work. Firstly, the case study approach used here extends the range of applicability of the concept of anomaly from general anomalies to those that are specifically concerned with human automation interaction. In other words, the focus on the historical trajectory seems a viable lens for producing constructive insights into the management of automation side effects.

Secondly, in Vaughan's work the anomaly trajectory was primarily exploited in order to understand how deviations from an accepted norm become increasingly routinised. In this research, the same focus shed light not only on the dynamics of the interpretation and framing of technology and its side effects, but also on the essential traits of successful and less successful organisational responses to such issues (framework components OP2 and OP3). In other words, while the use of the anomaly trajectory in Vaughan's original work was instrumental especially in illuminating processes of drift into failure, in this work the same concept was instrumental in understanding other areas of organisations that, as earlier anticipated, seem to be more relevant to the management of technology-



related anomalies, i.e., HAI issues.

These methodological considerations are important for future researchers interested in understanding the organisational contributors to HAI issues; especially those HAI issues that can be directly involved in the occurrence of accidents.

As introduced earlier (§ 4.3.1), the official investigation of automation failure in safety-critical domains is not an institutionalised practice—at least not to the same extent as accident investigations are. In fact, there are no specific agencies with a government mandate to investigate the causal factors leading to the deployment and use of poor automation—i.e., automation that does not deliver the intended benefit, automation that has to be removed from operation, or automation that introduces new types of errors. This is of course not to say that automation is not addressed in official investigations of transportation accidents: indeed, it is, but mainly as a contributory condition to the overall accident, and not as an issue in itself. However, due to the increasing pervasiveness of automation in safety-critical systems, it is desirable to expect that the importance of this type of investigations will increase in future. The more automation becomes an essential component of these systems, the higher the expectation of the educated public to have adequate governance structures in place to control it.

### 8.4. PRAGMATIC CONSIDERATIONS

---

The OPHAI categories can be turned into a checklist and/or questionnaire items for auditing purposes, so as to aid decision makers in safety-critical service providers in evaluating the adequacy of their organisational context to manage the HAI issues accompanying a given automated system. In particular, OPHAI can be helpful for:

1. **Periodic reviews of existing automation implementations.** The framework can be used as an aid in periodic reviews, to check how well the organisation is managing the automation side effects of an implemented system. For instance, if the automation comes to exhibit a poor fit with the intended users, i.e., if things start to get wrong, the framework can be used to explain what is going wrong and how—in particular, how automation unintended consequences are (mis) managed, and what corrective actions can be taken. For instance, during a periodic review, it may emerge that personnel hold ambiguous or conflicting views about a new system's operational role, or that they have a poor understanding of the system's HAI issues. Hence, senior management, once made aware of the problem, can be invited to direct more effort towards the understanding of the system's actual role, and the definition of internal

policies that clarify it. In summary, by checking the framework's components against the situation at hand, safety and human factors specialists can support project managers and programme leaders in formulating hypotheses about why and how the organisation is unable to manage HAI issues. This use of the framework is functional to ensure the monitoring of automation implementation quality past the O-date.

2. **Retrospective reviews of past (failed) automation implementations.** The framework can also support reviews of past (failed or successful) implementation projects. Analogously to periodic review of current automation implementations, the framework can stimulate project managers and programme leaders to consider the adequacy of the organisational conditions and managerial arrangements in place during the implementation of the programme in question.
3. **Comparative/benchmarking studies.** The framework can be used by policy makers, standardisation bodies, and regulators for collecting data about and compare the potential for controlling HAI issues across different service provider organisations. As the comparison between Alphasky and Deltasky suggests, different organisations may develop their potential to control automation's side effect at different points in time. (While at the time of data collection Alphasky was considered to be a successful case, Deltasky was found to be struggling with MSAW implementation.) Therefore, the framework could be functional to monitor the existence of potential asymmetries in the implementation quality of the same system across different service providers. Such asymmetries, if found to be relevant, can inform remedial interventions such as technology transfer programmes, to ensure that the knowledge necessary to control certain HAI issues is effectively transferred from best practice providers to less successful ones.
4. **Enhancement of (safety) reporting schemes.** The framework can be used to enhance existing (safety) reporting frameworks, such as the NASA Aviation Safety Reporting System database, to enable these to capture also organisational factors that are relevant to the management of HAI issues. These frameworks do not necessarily consider organisational precursors to HAI issues. The OPHAI framework can enhance these kinds of frameworks as its categories can provide a consistent format for collecting narrative reports about organisational precursors specific to HAI issues. The use of a consistent data format would permit organisations to store, communicate, aggregate and compare these reports in order to identify organisational weaknesses and deficiencies in the management and oversight of safety-critical automation. Subsequently, guidance for safety-critical automation improvement at local and

international level can be developed. Note that the identification of such organisational precursors would be mainly based on qualitative narratives, as these precursors, indeed, do not lend themselves to quantification, automatic reporting, and statistical comparison as for more operation near types of safety occurrences—i.e., CFITs, mid-air collisions, Loss of Control in Flight, as well as the immediate precursors conditions to these occurrences. However, the capture of the OPHAII's organisational precursors in the context of a broader reporting scheme is justified by the fact that, although not quantifiable, these precursors can lead to HAI issues that can be themselves a precursor to incidents and disasters. Hence, it is reasonable to track these precursors in order to promote the understanding of organisational weaknesses in the management of automation.

These uses of the framework are consistent with current standard risk management approaches, such as the ICAO's Safety Management System (SMS) (ICAO, 2013) and the ISO 31000 (ISO, 2009). These standard approaches provide guidance to organisations on how to establish formal frameworks for managing risk proactively. They include measures aimed at promoting the systematic and proactive identification and reporting of hazards and risks; continuous risk monitoring; process improvement; and timely mitigation of reported issues. However, both standards have a general purpose orientation in that they have been devised to manage different types of risks. Moreover, the ISO 31000 applies to different types of industries, not necessarily safety critical. Thus, the OPHAII's uses highlighted above can benefit those safety critical service provider organisations that want to include the handling of automation side effects in the scope of their ICAO SMS or ISO 31000. OPHAII can add to the current measures included under the two approaches as an additional supporting measure specifically tailored for the handling of HAI issues.

Having outlined the potential uses of OPHAII, it is important to identify what are the distinctive aspects of the framework that support these uses. These aspects are determined by means of a comparison with other pragmatic approaches found in the literature that were developed to address the same problem, i.e., HAI issues, or that include the consideration of organisational factors in their scope, although for different purposes, namely safety occurrence investigation and healthcare technology analysis.

### **8.4.1. Comparison with frameworks for handling HAI issues**

The direct relevance of OPHAII to the handling of HAI issues call for a comparison with approaches that have been developed to address the same problem. Bligård (2012) has

proposed a predictive framework for identifying the HAI issues accompanying the introduction of novel technology. While being useful in its own right to identify (and anticipate) possible HAI issues, such as error inducing conditions in the interface design, Bligård's framework focuses on the HAI level only; therefore, it does not consider the organisational precursors to HAI issues that are addressed, instead, by OPHAI.

Similar considerations apply in relation to the Interactive Socio-Technical Analysis (ISTA) (Harrison, Koppel, & Bar-Lev, 2007). Developed in the healthcare domain and included as a recommended method in the RAND's *Guide to Reducing Unintended Consequences of Electronic Health Records* (Jones et al., 2011), this framework provides a structured approach for investigating the root causes of adverse medical events induced by medical devices. More specifically, the approach advocates the construction of a timeline of the sequence of events that culminated with the undesirable event—e.g., a device administering an excessive dosage of a drug and consequential harm to patient. From here, the analyst can identify root causes that promoted the reconstructed event sequence. Such causes lie in the interaction of the medical device with the user and the surrounding work environment, and they may include factors such as workload increase following the introduction of the device, placement of the device in busy areas exposed to many potential distractions, or lack of a policy that mandates double-checks for risky medications (Jones et al., 2011).

In promoting the identification of these kinds of root causes, ISTA has the merit of avoiding reducing explanations of adverse medical events to medication error. However, the framework stops the investigation at the level of the sharp end; it does not generate questions about the broader organisational precursors at the blunt end that may have initially led to the deployment and operation of problematic automations in the first place. This is the level of analysis of direct concern to OPHAI.

Other relevant frameworks against which to compare OPHAI include safety occurrence investigation methods and automation implementation frameworks that incorporate predefined sets of organisational factors. These are addressed in the following two sections.

### **8.4.2. Comparison with comparable safety investigation methods**

Two notable example of safety investigation methods are the Safety Occurrence Analysis Methodology (SOAM) (Licu, Cioran, Hayward, & Lowe, 2007), which was developed for the investigation of safety occurrences in the air traffic management domain; and the

human factors framework for the investigation of adverse medical events, developed by Henriksen et al. for the healthcare domain (Henriksen, Dayton, Keyes, Carayon, & Hughes, 2008). Both models are based on Reason's Swiss Cheese model of accident causation (Reason, 1997). Therefore, although variations in their respective classification schemes exist, they both prescribe the considerations of (1) factors proximal to the event, e.g., barriers and human involvement; (2) contextual factors, e.g., the physical and technological aspects of the work environment that predisposed the conditions for the accident to unfold; and (3) latent conditions, e.g., organisational and management factors. The focus on these factors results in a comprehensive picture of the precursors to the safety occurrence under investigation. These occurrences may be a CFIT or a loss of separation incident for SOAM; or a medication or diagnosis mishap for the Henriksen et al.'s model. The precursors considered also include automation's contribution to the adverse event in question. However, although the two models identify automation's role in accident development, neither SOAM nor Henriksen et al.'s framework has technology or automation side effects as their main foci. Thus, the organisational factors they consider are essentially relevant to explain the (erroneous) behaviour of practitioners at the organisation sharp end. They do not identify possible organisational roots of HAI issues as OPHAI does.

### **8.4.3. Comparison with socio-technical frameworks of healthcare IT implementation**

Socio-technical frameworks of healthcare IT (HIT) implementation consist of multidimensional models intended to provide a systemic view of the multiple factors and interactions that must be considered when introducing IT in healthcare settings. Two notable examples of these models include Sittig and Singh's eight-dimensional socio-technical framework for studying the effectiveness and safety of HIT implementations (Sittig & Singh, 2010); and Cullen et al.'s Contextual Implementation Model (CIM) (Callen, Braithwaite, & Westbrook, 2008). These models promote the identification of problems that may arise when introducing such technology, which can then be addressed in order to increase the safety and efficiency—i.e., the success—of implementation. Two points distinguish OPHAI from these models.

The first point concerns the use of the same term—"organisational factor"—to denote a different kind of factors than those addressed by OPHAI. This remark applies mostly to Sittig and Singh's framework (2010). The framework includes the following dimensions: (1) hardware and software computing infrastructure; (2) clinical content; (3) human

computer interface; (4) people; (5) workflow and communication; (5) internal organisational policies, procedures and culture; (6) external rules, regulations and pressures; (8) system measurement and monitoring. The dimension which incorporates organisational factors is (5), internal organisational policies, procedures and culture. An example of the sub items of this dimension is provided by Meek et al., who used Sittig and Singh's framework to understand electronic health record implementations found in 12 UK National Health Care hospitals (Meeks, Takian, Sittig, Singh, & Barber, 2014). The study was based on secondary analysis of interview data previously collected. In addition to capturing a genuine organisational precursor, namely the availability of an adequate IT budget to support ongoing IT requirements, other organisational factors identified included data confidentiality issues and the increased risk of incorrect selection posed by having multiple record numbers per patient. However, from an OPHAI perspective, this kind of issues represent the outcome of specific organisational processes and conditions, rather than organisational factors per se, as OPHAI focuses on organisational factors at the blunt end.

The second point that distinguishes these models from OPHAI is the absence of a category of organisational precursors related to the interaction between the service provider organisation and the software manufacturer. In OPHAI, this interaction is addressed by OP3. Both Sittig and Singh's framework and CIM do not consider this kind of interaction, as they focus on either (i) organisational factors internal to the organisation, such the organisational context and the local unit context (CIM), or the internal policies, procedures and culture (Sittig & Singh's framework); or on (ii) factors external to the organisation, i.e., factors that pertain to the broader institutional and economic environment in which the organisation operates. However, while maintaining this focus they do not address inter-organisational factors that occur at the organisational boundaries of the organisation.

### **8.5. GENERALISABILITY**

---

OPHAI was developed in a very specific context: the organisational implementation and improvement of an automated system from the air traffic management domain, the MSAW, and the management of a specific MSAW-related HAI, the nuisance alert problem. Study 3 concluded that the framework seems to be relevant to the management of safety nets in general, rather than to the MSAW only. This focus calls into question the framework's generalisability to other contexts. Upon elaborating on this question, it should be noted that even a restricted applicability of the framework to the contexts of

safety nets and MSAW implementations is valuable from a societal perspective. These systems are key safety defences of the current ATM system; therefore, developing measures that can improve their management means contributing to improve public safety.

That said, it is reasonable to believe that OPHAI may be relevant for safety and human factors practitioners and programme managers involved in the definition and management of automated systems in safety-critical service provider organisations. In these organisations, there are at least two dynamics that may affect sensitivity to HAI issues. The first is the constant exposure to the indeterminacy of the technologies these organisations deploy. The intricacies of actual technology usage and the situated meanings that sharp end operators can project over the technology in use are not immediately visible or understandable for those decision makers and stakeholders at the blunt end, as the latter lack proximity to the operational environment. In safety-critical service provider organisations there is a constant gap between the views of the *technology-in-use* versus the view of the *technology-as-imagined*<sup>8</sup>. For stakeholders at the blunt end, the complexity of the technology in use will never be as apparent as it is for sharp end practitioners. Unless adequate resources are dedicated to address this problem, stakeholders at the blunt end will almost always run the risk of developing overly simplistic assumptions about the technology and the accompanying HAIs.

This is especially the case considering that safety-critical service providers are technology-intensive organisations. The continual demands for cost-effectiveness and safety, combined with the belief that technology is the answer to most operational problems, result in the constant deployment of novel automated tools in the operational environment. This increases the number of the coupling and interactions between components. As a result, new layers of technology-related complexity build on each other, and it remains difficult for decision makers and stakeholders outside to the operation room to keep track of how new technologies are actually used.

The second dynamic that may affect sensitivity to HAI issues is the fact that, in safety-critical industries, human-centred development does not enjoy the same importance as non-safety-critical industries, such as consumer electronics (Boy, 2012). For consumer

---

<sup>8</sup> The notions of *technology-as-imagined* and *technology-as-used* discussed here are an extension of the known notions of *work-as-imagined* and *work-as-done*, which have been highlighted in the OS area (Hollnagel, 2012a; Nemeth, 2008; Dekker, 2007).

electronics manufacturers the ability to understand in detail the end-user's experience is a key factor for achieving and sustaining their competitive advantage. In other words, investments in user experience research usually generate a return in the form of increased sales. This is not necessarily the case for software manufacturers developing technology for safety-critical domains. In these domains automation development is known to be driven by technology centred, not human centred, development processes (Boy, 2012).

For safety-critical service providers the development of human-centred technology is not necessarily at the core of the organisational mission. Automation is not an end in itself, but a means of improving the quality, safety, and efficiency of the safety-critical services they provide—e.g., the provision of air traffic control services to airline company; the provision of air, maritime, and railways transportation to passengers; the provision of healthcare services to the population (e.g., Blumenthal, 2009). It is around the provision of such services that the organisation has developed its main structure of roles and responsibilities.

Therefore, in this context, the importance of human centred design is subsumed to other organisational and business performance areas which are more directly linked to the core activities of the organisation—the activities that are more financially relevant for the organisation. These activities are not the development of technology, but the delivery of a safety-critical service, usually to the public, under conditions of scarcity. Even when the organisation dedicates resources to novel technology development, these resources must be relatively limited. The majority of them is normally directed towards the administrative and operational activities necessary to sustain the services that generate (the majority of) company revenues. Thus, despite the continuous drive for the modernization of equipment and infrastructure, the integration of human centred design processes is not important here as it is for other industries.

Essentially, the conditions mentioned above, i.e., the indeterminacy of technology combined with the devaluing of human centred automation principles, entail the constant risk for safety-critical service providers of downplaying the complexity intrinsic in the use and operation of automation.

### **8.6. LIMITATIONS & FUTURE WORK**

---

The research managed to consider the perspectives of various stakeholders involved in the implementation and improvement of the MSAW. However, one limitation was the limited consideration of manufacturer representatives and senior management. Future



studies could collect more data from manufacturers, so as to provide deeper insights into the third category of the framework, which addresses the relationships between the service provider and the manufacturer. In particular, it would be useful to clarify the range of strategies that safety-critical service provider can exploit to ensure that the supplied system fits the specific operational environment. The limited involvement of senior management was mitigated in the present research by using documentary sources. Future studies can increase the focus on this stakeholder group in order to shed further light into the actual rationales behind the introduction and use of technology.

Furthermore, future studies may give further consideration to the perspectives of stakeholders external to the service provider organisation. Regarding such stakeholder group, this research considered the perspectives of the accident investigation body (Study 1), and those of a standard developer's agency (Studies 2 and 3). Future studies may further include the views of national and international regulators. Because they had a limited involvement in the setup of the studied alarm system, their views were minimally considered in this study (one EASA regulator was considered in study 2 and 3). However, it is likely that their involvement in the management of automation will increase as insights and lessons learned into the management of automation in safety-critical service providers, such as those provided by this study, become more available, and policies and regulations can be developed.

The research detailed here is centred on the MSAW and the problem of nuisance alerts. Although an argument has been put forward to justify the relevance of the OPHAI framework also to other contexts, future research may investigate the management of other HAI issues, such as problems of automation surprises, expertise degradation, data overload, related to other kinds of automated systems, e.g., information displays and decision support systems. Future studies may also consider automated systems implemented and managed by safety-critical service provider organisations other than ANSPs, such healthcare providers, airline operators, railways operators etc. Applying the framework in different contexts would provide a further point of triangulation to refine and increase the framework's generalisability.

Also, it can be noted that while the framework lists a set of conditions that are relevant to handling HAI issues, it provides limited insight into how these conditions can be successfully established over time and maintained in safety-critical service providers. Therefore, it is important to investigate further (i) how safety-critical service providers organisations actually develop (or fail to develop) appropriate assumptions about the

usage of their automated systems and the accompanying HAI issues; (ii) how they develop appropriate HAI issues handling capabilities; and (iii) how they establish control over implementation at the service provider software provider boundary.

Finally, future work may also investigate OPHAI's actual usefulness. As discussed earlier (§ 8.4), the framework has a range of potential pragmatic applications. Future work could actually develop these pragmatic applications and evaluate the added value that they may bring to specific organisations. At a minimum, this value can be investigated both retrospectively and comparatively. A retrospective analysis would look at the novel insights that the framework could produce in a case of a past situation where a specific set of HAI issues was poorly handled. A comparative analysis would require implementing the framework in a defined automation implementation and improvement context, in order to see, in this context, what insights into the organisational precursors to HAI issues such an application can deliver in comparison to a baseline context—i.e., a comparable context in which the framework is not used.

In concluding this section, it is important to remind that dedicated strategies have been implemented both at the level of the overall research (§ 4.3.4) and the individual studies (§ 4.6) in order to minimise the sources of bias and improve the validity of the study.

## **8.7. CONCLUSIONS**

---

The main contributions to knowledge delivered by this research are theoretical, pragmatic, and methodological. These are described in the following three sections.

### **8.7.1. Theoretical contribution**

The OPHAI framework contributes to addressing the research gap between current theoretical perspectives on HAI issues, the HCI (§ 2.2) and the system lifecycle (§ 2.3) perspectives, and the “enlarged” OS perspective (chapter 3), by identifying three classes of organisational precursors to HAI issues that can be found in the context of implementation and use of automation. Such precursors include: the organisational assumptions driving automation adoption and improvement; the availability of dedicated organisational capability for handling HAI issues; and the control over implementation quality established at the service provider software manufacturer boundary. The identification of these organisational precursors, which is supported by the findings of the research’s three empirical studies (chapters 5, 6, and 7), first, advances knowledge of how organisations operating in safety-critical domains actually manage the HAI issues accompanying the automated systems these organisations deploy and operate. Second, it supports the view of HAI issues as a symptom of deeper organisational problems, rather than solely as a problem of poor human machine interface design, or poor human factors and safety assurance.

### **8.7.2. Pragmatic contribution**

OPHAI framework’s pragmatic value lies in the fact that its categories identify traits of an organisational context that can be monitored in order to assess the adequacy of such a context to handle HAI issues. As discussed earlier (§ 8.4), the framework can be translated into a checklist or auditing tool that safety-critical service provider organisations can use for periodic or retrospective assessment of the adequacy of their organisational context to handle HAI issues. In addition, the framework can be used by policy makers, regulators or standard developers to carry out comparative assessments of the potential for handling HAI issues across different safety-critical service providers found in the same industry. A final possible use is the enhancement of existing reporting schemes, so as to allow the systematic capture of organisational weaknesses in the handling of automation. These pragmatic uses of the framework are important as they promote the identification of corrective actions that are targeted at the organisational level, and that are complementary to those identified by the HCI and the system lifecycle perspectives.

Another pragmatic contribution can be identified. The first two studies of the research represent a contribution in their own right, as they provide deep insights into the organisational experiences and conditions relevant to the implementation and improvement of the MSAW system and the accompanying main HAI issue. The ability of practitioners to act in a pragmatic area, such as system safety and human factors, depends on their knowledge of a repertoire of relevant cases (Johansson, 2003). The cases presented in Studies 1 and 2 can be considered as contribution as they contribute to the development of a repertoire of case studies regarding the management of safety-critical automation.

### **8.7.3. Methodological contribution**

This work makes three methodological contributions:

- The first methodological contribution lies in the use of an original unit of analysis: the organisational trajectory of an HAI issue. While building on foundational work from the OS area (§ 4.3.2), the present research has shown, by means of Studies 1 and 2, the insights that such a unit can provide, thus evincing the value this unit can deliver when used to inquire into the organisational and managerial sources of automation side effects in safety-critical domains (§ 8.3.2).
- The second methodological contribution lies in having shown the theoretical insights that can be gained by using a type of data not normally used in the study of technology or in the OS area. The present research shows the valuable insights this type of data can offer to studies with an interest in the management of technology and safety (§ 8.3.1).
- The third methodological contribution lies in the definition and use of an overall qualitative research strategy that, although based on the single qualitative case study approach, which is frequently found in the OS area, also departed from this basic approach (as described under § 4.3) in order to overcome its main limitations.

These methodological aspects of the work can be classified as contributions because they provide examples of methodological strategies and choices that have worked in this research and can therefore be replicated in future research on the interaction between technology and safety-critical organisations.

# References

- AAIB (2004). *Final Report No. 1793 by the Aircraft Accident Investigation Bureau concerning the accident to the aircraft AVRO 146-RJ100, HB-IXM, operated by Crossair under flight number CRX 3597, on 24 November 2001 near Bassersdorf/ZH* (Accident Investigation Report). Aircraft Accident Investigation Bureau, Berne.
- Bainbridge, L. (1983). Ironies of automation. *Automatica*, 19(6), 775–779.
- Balka, E., Doyle-Waters, M., Lecznarowicz, D., & FitzGerald, J. M. (2007). Technology, governance and patient safety: Systems issues in technology and patient safety. *International Journal of Medical Informatics*, 76, 35–47.
- Balka, E., & Kahnamoui, N. (2004). Technology Trouble? Talk to Us: Findings from an Ethnographic Field Study. In *Proceedings of the Eighth Conference on Participatory Design: Artful Integration: Interweaving Media, Materials and Practices - Volume 1* (pp. 224–234). New York City, NY, USA: ACM.
- Bannon, L. (2000). Understanding common information spaces in CSCW. In *Workshop on Cooperative Organisation of Common Information Spaces*, Technical University of Denmark.
- Bannon, L., & Bødker, S. (1997). Constructing common information spaces. In *Proceedings of the Fifth European Conference on Computer Supported Cooperative Work* (pp. 81–96). Springer Netherlands.
- Bardram, J. E., & Bossen, C. (2005). A web of coordinative artifacts: collaborative work at a hospital ward. In *Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work* (pp. 168–176). ACM.
- BEA. (2012). *Final Report on the accident on 1st June 2009 to the Airbus A330-203 registered F-GZCP operated by Air France, flight AF 447 Rio de Janeiro - Paris* (Accident Investigation Report). Bureau d'Enquêtes et d'Analyses pour la sécurité de l'aviation civil.
- Beck, U. (2000). Risk society revisited: theory, politics and research programmes. *The Risk Society and beyond: Critical Issues for Social Theory*, 211–29.
- Benbasat, I., Goldstein, D. K., & Mead, M. (1987). The case research strategy in studies of information systems. *MIS Quarterly*, 369–386.
- Bhaskar, R. (1979). *The possibility of naturalism: A philosophical critique of the contemporary human sciences*. Brighton, UK: Harvester Press.

- Blau, Peter M. (1964) *Exchange and Power in Social Life*. New York: John Wiley.
- Blaikie, N. (2009). *Designing Social Research* (2nd edition). Cambridge, UK: Polity Press.
- Bligård, L.-O. (2012). *Predicting mismatches in user-artefact interaction - Development of an analytical methodology to support design work* (PhD Thesis). Department of Product and Production Development, Chalmers University, Gothenburg, Sweden.
- Blumenthal, D. (2009). Stimulating the Adoption of Health Information Technology. *New England Journal of Medicine*, 360(15), 1477–1479.
- Bonneau, C. (2013). Contradictions and their concrete manifestations: an activity-theoretical analysis of the intra-organizational co-configuration of open source software. EGOS.
- Bourrier, M. (2002). Bridging Research and Practice: The Challenge of 'Normal Operations' Studies. *Journal of contingencies and crisis management*, 10, 173-180.
- Bourrier, M. (2011). *The Legacy of the Theory of High Reliability Organizations: An Ethnographic Endeavour*. Genève: Université de Genève.
- Boy, G. (2012). *Orchestrating human-centered design*. Springer Science & Business Media.
- Burrell, G., & Morgan, G. (1979). *Sociological paradigms and organisational analysis* (Vol. 248). London: Heinemann.
- Busby, J. S., & Bennett, S. A. (2007). Loss of defensive capacity in protective operations: the implications of the Überlingen and Linate disasters. *Journal of Risk Research*, 10(1), 3–27.
- Callen, J. L., Braithwaite, J., & Westbrook, J. I. (2008). Contextual Implementation Model: A Framework for Assisting Clinical Information System Implementations. *Journal of the American Medical Informatics Association : JAMIA*, 15(2), 255–262.
- Campbell, E. M., Sittig, D. F., Ash, J. S., Guappone, K. P., & Dykstra, R. H. (2006). Types of Unintended Consequences Related to Computerized Provider Order Entry. *Journal of the American Medical Informatics Association*, 13(5), 547–556.
- Capra, F. (1997). *La rete della vita* (Vol. 38). Milano: Rizzoli.
- Cardosi, K. (1998). Human Factors Lessons Learned in the Design and Implementation of Air Traffic Control Systems. *The Controller, First Quarter*, 11–15.
- Carlisle, L. A. (2000). FAA v. the NTSB: Now That Congress Has Addressed the Federal Aviation Administration's 'Dual Mandate', Has the FAA Begun Living up to Its Amended Purpose of Making Air Travel Safer, or Is the National Transportation Safety Board Still Doing Its Job Alone? *J. Air L. & Com.*, 66, 741.
- Catino, M. (2002). *Da Chernobyl a Linate. Incidenti tecnologici o errori organizzativi?* Roma: Carrocci.
- Catino, M. (2013). *Organizational myopia: Problems of rationality and foresight in organizations*. Cambridge: Cambridge University Press.
- Checkland, P. (1999). *Systems thinking, systems practice: includes a 30-year retrospective*. Chichester, UK: John Wiley & Sons.
- Choudhury, V., & Sabherwal, R. (2003). Portfolios of control in outsourced software development projects. *Information Systems Research*, 14(3), 291–314.
- Cicognani, E. (2002). L'approccio qualitativo della Grounded Theory in psicologia sociale: Potenzialità, ambiti di applicazione e limiti. *Metodi qualitativi in psicologia sociale*, 43-59.
- Cook, R., Nemeth, C., & Dekker, S. (2008). What went wrong at the Beatson Oncology

- Centre. In E. Hollnagel, C. Nemeth, & S. Dekker (Eds.), *Resilience engineering perspectives* (Vol. 1: Remaining Sensitive to the Possibility of Failure, pp. 225–236). Aldershot, UK: Ashgate Publishing, Ltd.
- Cook, R., & Rasmussen, J. (2005). 'Going solid': a model of system dynamics and consequences for patient safety. *Quality and Safety in Health Care, 14*(2), 130–134.
- Corbetta, P. (2003). *La ricerca sociale: metodologia e tecniche*. Il Mulino.
- Cornelissen, J. P. (2006). Making sense of theory construction: Metaphor and disciplined imagination. *Organization Studies, 27*(11), 1579–1597.
- Crotty, M. J. (1998). *The Foundations of Social Research: Meaning and Perspective in the Research Process*. London, UK: SAGE.
- Dalcher, D. (2004). Stories and histories: Case study research (and beyond) in information systems failures. *The Handbook of Information Systems Research, 305–322*.
- Dalcher, D. (2009). Making sense of IS failures. *Encyclopedia of Information Science and Technology, 5, 2476–83*.
- Danko, M. (2010). FAA Ignores NTSB Safety Recommendations: Aviation Law Monitor. Retrieved August 25, 2012, from [www.aviationlawmonitor.com/2010/10/articles/ntsb/faa-ignores-ntsb-safety-recommendations/](http://www.aviationlawmonitor.com/2010/10/articles/ntsb/faa-ignores-ntsb-safety-recommendations/).
- Degani, A. (2004). *Taming HAL: Designing interfaces beyond 2001*. New York: Palgrave Macmillan.
- Dekker, S. (2007). Resilience engineering: Chronicling the emergence of confused consensus. In E. Hollnagel, D. D. Woods, & N. Leveson (Eds.), *Resilience Engineering: Concepts and Precepts*. Ashgate Publishing, Ltd.
- Dekker, S. (2011). *Patient Safety: A Human Factors Approach*. Boca Raton: Crc Press.
- Dekker, S. (2012). *Drift into Failure: From Hunting Broken Components to Understanding Complex Systems*. Ashgate Publishing, Ltd.
- Demchak, C. C. (1991). *Military Organizations, Complex Machines: Modernization in the U.S. Armed Services*. Ithaca, N.Y.: Cornell University Press.
- Dennis, L. (2005). *Scenario Network Planning: The development of a Methodology for Social Inquiry* (PhD Thesis). Division of Business and Enterprise. University of South Australia, Adelaide.
- DOT/FAA. (1995). Order 1220.2F: FAA Procedures for handling National Transportation Safety Board Recommendations. Retrieved May 6, 2009, from [www.faa.gov/documentLibrary/media/Order/ND/1220.2F.pdf](http://www.faa.gov/documentLibrary/media/Order/ND/1220.2F.pdf).
- Dubois, A., & Gadde, L.-E. (2014). 'Systematic combining'—A decade later. *Journal of Business Research, 67*(6), 1277–1284.
- Dutch Safety Board. (2010). *Crashed during approach, Boeing 737-800, near Amsterdam Schiphol Airport, 25 February 2009* (Accident Investigation Report). The Hague.
- Edmondson, A., Ferlins, E., Feldman, L., & Bohmer, R. (2005). The Recovery Window: Organizational Learning Following Ambiguous Threats. In M. Farjoun & W. H. Starbuck (Eds.), *Organization at the Limit: Lessons from the Columbia Disaster*. Oxford, UK: Blackwell.
- Engestrom, Y. (2000). Activity theory as a framework for analyzing and redesigning work. *Ergonomics, 43*(7), 960–974.
- Ericson, C. A. (2005). *Hazard analysis techniques for system safety*. Hoboken, New Jersey, USA: John Wiley & Sons.
- EUROCONTROL. (No date). Safety Nets. Retrieved April 2, 2015, from

- www.eurocontrol.int/safety-nets.
- EUROCONTROL. (2001). *ESARR 4: Risk Assessment and Mitigation in ATM Edition 1.0*. Brussels.
- FAA. (No date). FAA Historical Chronology, 1926-1996. Retrieved March 15, 2009, from [www.faa.gov/about/media/b-chron.pdf](http://www.faa.gov/about/media/b-chron.pdf).
- FAA. (1981, December 21). Response letter to NTSB safety recommendation A-81-134. Retrieved September 5, 2014, from [www.nts.gov/SafetyRecs/Private/QueryPage.aspx](http://www.nts.gov/SafetyRecs/Private/QueryPage.aspx).
- FAA. (1985a, January 18). Response letter to NTSB safety recommendation A-84-83. Retrieved September 5, 2014, from [www.nts.gov/SafetyRecs/Private/QueryPage.aspx](http://www.nts.gov/SafetyRecs/Private/QueryPage.aspx).
- FAA. (1985b, January 18). Response letter to NTSB safety recommendation A-84-84. Retrieved September 5, 2014, from [www.nts.gov/SafetyRecs/Private/QueryPage.aspx](http://www.nts.gov/SafetyRecs/Private/QueryPage.aspx).
- FAA. (1985c, August 26). Response letter to NTSB letter dated 1 July 1985 [on A-84-83]. Retrieved September 5, 2014, from [www.nts.gov/SafetyRecs/Private/QueryPage.aspx](http://www.nts.gov/SafetyRecs/Private/QueryPage.aspx).
- FAA. (1985d, August 26). Response letter to NTSB letter dated 1 July 1985 [on A-84-84]. Retrieved September 5, 2014, from [www.nts.gov/SafetyRecs/Private/QueryPage.aspx](http://www.nts.gov/SafetyRecs/Private/QueryPage.aspx).
- FAA. (1991, January 18). Response letter to NTSB safety recommendation A-90-161. Retrieved September 5, 2014, from [www.nts.gov/SafetyRecs/Private/QueryPage.aspx](http://www.nts.gov/SafetyRecs/Private/QueryPage.aspx).
- Farjoun, M. (2005). Organizational Learning and Action in the Midst of Safety Drift: Revisiting the Space Shuttle Program's Recent History. In M. Farjoun & W. H. Starbuck (Eds.), *Organization at the Limit: Lessons from the Columbia Disaster* (pp. 60–80). Oxford, UK: Blackwell.
- Fields, B., Amaldi, P., & Tassi, A. (2005). Representing collaborative work: the airport as common information space. *Cognition, Technology & Work*, 7(2), 119–133.
- Flick, U. (2009). *An introduction to qualitative research*. London, UK: Sage.
- Flin, R., O'Connor, P., & Crichton, M. (2008). Safety at the sharp end. *A guide to nontechnical skills*. Farnham, UK: Ashgate.
- Flynn, R. (2006). Health and risk. In *Beyond the Risk Society: Critical Reflections on Risk and Human Security* (pp. 77–95). Maidenhead: UK: Open University Press.
- Fortune, J., & Peters, G. (1995). *Learning from failure - The systems approach*. Chichester, UK: Wiley.
- Glaser, B. G., & Strauss, A. L. (2009). *The Discovery of Grounded Theory: Strategies for Qualitative Research*. Transaction Publishers.
- Grounds, C. B., & Ensing, A. R. (2000). Automation Distrust Impacts on Command and Control Decision Time. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 44(22), 855–858.
- Goffman, Erving (1961) *Asylums: Essays on the Social Situation of Mental Patients and Other Inmates*. New York: Anchor.
- Hall, P. G. (1982). *Great planning disasters*. Berkley, CA: University of California Press.
- Harris, D. (2012). *Human Performance on the Flight Deck*. Aldershot, UK: Ashgate.
- Harrison, M. I., Koppel, R., & Bar-Lev, S. (2007). Unintended consequences of information technologies in health care—an interactive sociotechnical analysis. *Journal of the American Medical Informatics Association*, 14(5), 542–549.



- Heath, C., & Luff, P. (1991). Collaborative activity and technological design: Task coordination in London Underground control rooms. In *Proceedings of the Second European Conference on Computer-Supported Cooperative Work ECSCW'91* (pp. 65–80). Amsterdam, The Netherlands: Springer.
- Henriksen, K., Dayton, E., Keyes, M. A., Carayon, P., & Hughes, R. (2008). Understanding Adverse Events: A Human Factors Framework. In R. G. Hughes (Ed.), *Patient Safety and Quality: An Evidence-Based Handbook for Nurses*. Rockville (Maryland): Agency for Healthcare Research and Quality (USA).
- HFI DTC. (2007). *The MOD HFI Process Handbook*. MBDA UK Ltd.
- Hirschheim, R. (1985). Information systems epistemology: An historical perspective. *Research methods in information systems*. In E. Mumford, R. Hirschheim, G. Fitzgerald, & A. T. Wood-Harper, A. T. (Eds.), *Research methods in information systems*. North-Holland Publishing Co.
- Hirschman, Albert O. (1970) Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States. Cambridge, MA: Harvard University Press.
- Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction*, 7(2), 174–196.
- Hollnagel, E. (2004). Barriers and accident prevention. Aldershot, UK: Ashgate.
- Hollnagel, E., Woods, D. D., & Leveson, N. (2006). *Resilience engineering: Concepts and precepts*. Ashgate Publishing, Ltd.
- Hollnagel, E. (2012a). Resilience engineering and the systemic view of safety at work: Why work-as-done is not the same as work-as-imagined. In *Bericht zum 58. Kongress der Gesell - schaft für Arbeitswissenschaft vom 22 bis 24 Februar 2012* (pp. 19–24). Dortmund, Germany: Gfa-Press.
- Hollnagel, E. (2012b). *The ETTO Principle: Efficiency-Thoroughness Trade-Off*. Farnham, UK: Ashgate.
- Hollnagel, E., & Woods, D. D. (2005). *Joint Cognitive Systems: Foundations of Cognitive Systems Engineering*. New York: CRC Press.
- Hollnagel, E. (2014). Safety-I and Safety-II: The Past and Future of Safety Management. Ashgate Publishing, Ltd.
- Hopkins, A. (2014). Issues in safety science. *Safety Science*, 67, 6–14.
- Howell, R. (2011). NUISANCE ALERTS - More than just a numbers game. *EUROCONTROL NETALERT Newsletter 13*, (13), 5–6.
- Hughes, Everett C. (1984) *The Sociological Eye*. New Brunswick, NJ: Transaction Books.
- Humphreys, P., Kirwan, B., & Ternov, S. (2006). *A safe approach to transition: Key elements for transition success* (EEC Technical / Scientific Reports). EUROCONTROL.
- ICAO. (2013). *Doc 9859 Safety Management Manual (SMM) (3rd edition)*. Montréal, Canada: International Civil Aviation Organization.
- ISO. (2009). ISO 31000: 2009 Risk management--Principles and guidelines. International Organization for Standardization.
- ISO. (2010). ISO 9241-210: 2010 Ergonomics of human-system interaction -- Part 210: Human-centred design for interactive systems. International Organization for Standardization.
- Jones, S. S., Koppel, R., Ridgely, M. S., Palen, T. E., Wu, S.-Y., & Harrison, M. I. (2011). Guide to Reducing Unintended Consequences of Electronic Health Records.
- Johansson, R. (2003). Case study methodology. In *the International Conference on*

- Methodologies in Housing Research, Stockholm, Sweden 22–24 September 2003.*
- Keller, J. P., Diefes, R., Graham, K., Meyers, M., & Pelczarski, K. (2011). Why clinical alarms are a 'top ten' hazard: how you can help reduce the risk. *Biomedical Instrumentation & Technology, 45*(s1), 17–23.
- Kern, T., & Willcocks, L. (2001). Contract, Control and 'Presentation' in IT Outsourcing: Research in Thirteen UK Organizations. *Advanced Topics in Global Information Management, Volume 1*, 227.
- Kirsch, L. J., Sambamurthy, V., Ko, D.-G., & Purvis, R. L. (2002). Controlling information systems development projects: The view from the client. *Management Science, 48*(4), 484–498.
- Kirwan, B., & Ainsworth, L. K. (1992). *A guide to task analysis: the task analysis working group*. London, UK: Taylor & Francis.
- Kirwan, B. (2007). Safety informing design. *Safety Science, 45*(1–2), 155–197.
- Kirwan, B., Evans, A., Donohue, L., Kilner, A., Lamoureux, T., Atkinson, T., & MacKendrick, H. (1997). Human Factors in the ATM System Design Life Cycle. Presented at the FAA/Eurocontrol ATM R&D Seminar, Paris, France.
- Koppel, R., Metlay, J. P., Cohen, A., Abaluck, B., Localio, A. R., Kimmel, S. E., & Strom, B. L. (2005). Role of computerized physician order entry systems in facilitating medication errors. *The Journal of the American Medical Association, 293*(10), 1197–1203.
- Johnson, C. W., & de Almeida, I. M. (2008). An investigation into the loss of the Brazilian space programme's launch vehicle VLS-1 V03. *Safety Science, 46*(1), 38–53.
- Johnson, R. B. (1997). Examining the validity structure of qualitative research. *Education, 118*(2), 282.
- LaPorte, T.R. (1988). The United States air traffic system: Increasing reliability in the midst of rapid growth. In R. Mayntz and T. Hughes (Eds.), *The development of large scale technical systems* (pp. 215–244). Boulder: Westview Press.
- La Porte, T. R. (1996). High Reliability Organizations: Unlikely, Demanding and At Risk. *Journal of Contingencies and Crisis Management, 4*(2), 60–71.
- LaPorte, T. R., & Consolini, P. M. (1991). Working in practice but not in theory: theoretical challenges of 'high-reliability organizations'. *Journal of Public Administration Research and Theory, 19–48*.
- Leinwand, P., & Mainardi, C. R. (2010). *The Essential Advantage: How to Win with a Capabilities-Driven Strategy*. Boston, Massachusetts: Harvard Business Review Press.
- Le Coze, J. C. (2008). Disasters and organisations: From lessons learnt to theorising. *Safety Science, 46*(1), 132–149.
- Le Coze, J. C., & Pettersen, K. (2008). Is resilience engineering realist or constructivist? In *3<sup>rd</sup> Resilience Engineering Symposium* (pp. 175–184). École des Mines de Paris. Paris.
- Leonard-Barton, D. (1998). *Wellsprings of Knowledge: Building and Sustaining the Sources of Innovation*. Harvard Business Press.
- Leveson, N. (2012). *Engineering a Safer World: Systems Thinking Applied to Safety*. Cambridge, Massachusetts, USA: The MIT Press.
- Leveson, N., de Villepin, M., Srinivasan, J., Daouk, M., Neogi, N., Bachelder, E., ... Flynn, G. (2001). A Safety and Human-Centered Approach to Developing New Air Traffic Management Tools. In *4th USA/Europe Air Traffic Management R&D Seminar*.
- Licu, T., Cioran, F., Hayward, B., & Lowe, A. (2007). EUROCONTROL—Systemic Occurrence

- Analysis Methodology (SOAM)—A 'Reason'-based organisational methodology for analysing incidents and accidents. *Reliability Engineering & System Safety*, 92(9), 1162–1169.
- Loewenstein, J., & Ocasio, W. 2004. Vocabularies of organizing: linking language, culture, and cognition in organizations. Unpublished manuscript, Northwestern University.
- Lowe, C. (2008). A human factors perspective on safety management systems. In F. Redmill & T. Anderson (Eds.), *Improvements in System Safety* (pp. 139–153). London, UK: Springer.
- Marais, K., & Saleh, J. H. (2008). Conceptualizing and communicating organizational risk dynamics in the thoroughness–efficiency space. *Reliability Engineering & System Safety*, 93(11), 1710–1719.
- Marais, K., Saleh, J. H., & Leveson, N. G. (2007). Organizational risk dynamics in complex goal environments. In *Proceedings of ESREL*.
- March, J.G., & Simon, H.A. (1958). *Organizations*. New York: Wiley.
- Maxwell, J. (1992). Understanding and validity in qualitative research. *Harvard educational review*, 62(3), 279–301.
- Mayring, P. (2010). *Qualitative Inhaltsanalyse [Qualitative content analysis. Basics and techniques]* (11th edition). Weinheim, Germany: Beltz.
- Meeks, D. W., Takian, A., Sittig, D. F., Singh, H., & Barber, N. (2014). Exploring the sociotechnical intersection of patient safety and electronic health record implementation. *Journal of the American Medical Informatics Association*, 21(e1), e28–e34.
- Miles, M. B., & Huberman, A. M. (1994). *Qualitative Data Analysis: An Expanded Sourcebook(2nd edition)* (2nd edition). Thousand Oaks, California, USA: SAGE.
- Milliken, F., Theresa, L., & Bridewell-Mitchell, E. (2005). Barriers to the Interpretation and Diffusion of Information about Potential Problems in Organizations: Lessons from the Space Shuttle Columbia. In M. Farjoun & W. Starbuck (Eds.), *Organization at the Limit: Lessons from the Columbia Disaster* (pp. 246–266). Oxford, UK: Blackwell.
- Morgan, G. (1997). *Imaginization: New Mindsets for Seeing, Organizing and Managing*. Thousand Oaks, California, USA: Berrett-Koehler Publishers and SAGE.
- Morgan, G. (2006). *Images of Organization* (Updated edition). Thousand Oaks, California, USA: SAGE Publications.
- Naikar, N. (2013). *Work Domain Analysis: Concepts, Guidelines, and Cases*. Boca Raton, Florida, USA: CRC Press.
- Nebeker, J. R., Hoffman, J. M., Weir, C. R., Bennett, C. L., & Hurdle, J. F. (2005). High rates of adverse drug events in a highly computerized hospital. *Archives of Internal Medicine*, 165(10), 1111–1116.
- Nemeth, C. P. (2008). Resilience engineering: The birth of a notion. In E. Hollnagel, C. P. Nemeth, & S. Dekker (Eds.), *Resilience Engineering Perspectives: Remaining sensitive to the possibility of failure* (Vol. 1). Ashgate Publishing, Ltd.
- Norman, D. A. (1990). The problem with automation: inappropriate feedback and interaction, not 'over-automation'. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 327(1241), 585–593.
- NRI, 2009. NRI MORT User-s Manual (2nd Edition). Noordwijk Risk Initiative Foundation. The Netherland.
- NTSB. (1973). *Eastern Air Lines, Inc, L-1011, N310EA, Miami, Florida, December 29, 1972*

- (Aircraft Accident Report No. NTSB/AAR-73/14). Washington, DC, USA.
- NTSB. (1981). *Special Investigation Report, Aircraft Separation Incidents at Hartfield, Atlanta International Airport, Atlanta, Georgia, October 7, 1980* (No. PB82-123985). Washington, DC, USA.
- NTSB. (1984, August 13). NTSB Safety recommendation(s) A-84-82 through -84. Retrieved September 5, 2014, from [www.nts.gov/SafetyRecs/Private/QueryPage.aspx](http://www.nts.gov/SafetyRecs/Private/QueryPage.aspx).
- NTSB. (1985, July 1). Response letter to FAA letter dated 1/18/1985 [on A-84-84]. Retrieved September 5, 2014, from [www.nts.gov/safetyrecs/private/history.aspx?rec=A-84-084&addressee=FAA](http://www.nts.gov/safetyrecs/private/history.aspx?rec=A-84-084&addressee=FAA).
- NTSB. (1986, January 24). Response letter to FAA letter dated 8/26/1985 [on A-84-84]. Retrieved September 5, 2014, from [www.nts.gov/safetyrecs/private/history.aspx?rec=A-84-084&addressee=FAA](http://www.nts.gov/safetyrecs/private/history.aspx?rec=A-84-084&addressee=FAA).
- NTSB. (1990, October 29). NTSB Safety recommendation A-90-160 through -163. Retrieved September 5, 2014, from [www.nts.gov/SafetyRecs/Private/QueryPage.aspx](http://www.nts.gov/SafetyRecs/Private/QueryPage.aspx).
- NTSB. (2000). *Controlled Flight Into Terrain, Korean Air Flight 801, Boeing 747-300, HL7468, Nimitz Hill, Guam, August 6, 1997* (Accident Investigation Report No. NTSB/AAR-00/01). Washington, DC, USA.
- NTSB. (2004a). About NTSB. History and Mission. Retrieved May 12, 2009, from [www.nts.gov/Abt\\_NTSB/history.htm](http://www.nts.gov/Abt_NTSB/history.htm).
- NTSB. (2004b). Recommendations & Advocacy. Retrieved May 12, 2009, from [www.nts.gov/Recs/index](http://www.nts.gov/Recs/index).
- NTSB. (2014). *Descent Below Visual Glidepath and Impact With Seawall Asiana Airlines Flight 214, Boeing 777-200ER, HL7742, San Francisco, California* (No. NTSB/AAR-14/01). Washington, DC: National Transportation Safety Board.
- Ocasio, W. (2005). The opacity of risk: Language and the culture of safety in NASA's space shuttle program. In W. Starbuck & M. Farjoun (Eds.), *Organization at the limit: Lessons from the Columbia disaster*. Oxford, UK: Blackwell.
- Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 39(2), 230–253.
- Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 30(3), 286–297.
- Perin, C. (2006). *Shouldering Risks: The Culture of Control in the Nuclear Power Industry*. Princeton, New Jersey, USA: Princeton University Press.
- Partington, D. (2002). Grounded theory. In D. Partington (Ed.), *Essential skills for management research*, London, UK: Sage, pp 136-157.
- Perrow, C. (2011). *Normal Accidents: Living with High Risk Technologies*. Princeton University Press.
- Perry, S. J., Wears, R. L., & Cook, R. I. (2005). The role of automation in complex system failures. *Journal of Patient Safety*, 1(1), 56–61.
- Pettersen, K. (2008). The relationship between social reality and the physical – A causal nexus for operational safety and vulnerability. In Svedung, I., Enander, A. & Axelsson R. (Eds.), *Anthology on Learning from Accidents, an anthology based on thoughts and ideas from young research fellows*, Swedish Rescue Services Agency.
- Phillips, E., & Derek, P. (2010). *How to get a PhD: a handbook for students and their supervisors (5h edition)* (5 edition). Maidenhead, UK: Open University Press.

- Pidgeon, N., & O'Leary, M. (2000). Man-made disasters: why technology and organizations (sometimes) fail. *Safety Science*, 34(1), 15–30.
- Porte, T. L., & Consolini, P. (1998). Theoretical and operational challenges of 'high-reliability organizations': air-traffic control and aircraft carriers. *International Journal of Public Administration*, 21(6-8), 847–852.
- Prahalad, C. K., & Hamel, G. (1990). *The core competence of the corporation* (Vol. 1990). Boston, Massachusetts, USA: Harvard Business School Publishing Corporation.
- Rajkumar, A., & Blandford, A. (2012). Understanding infusion administration in the ICU through Distributed Cognition. *Journal of Biomedical Informatics*, 45(3), 580–590.
- Rasmussen, J. (1997). Risk management in a dynamic society: a modelling problem. *Safety Science*, 27(2–3), 183–213.
- Rasmussen, J., Pejtersen, A. M., & Goodstein, L. P. (1994). *Cognitive Systems Engineering* (1 edition). New York City, New York, USA: Wiley-Interscience.
- Rasmussen, J., & Svedung, I. (2000). *Proactive Risk Management in a Dynamic Society*. Karlstad, Sweden: Swedish Rescue Services Agency.
- Reason, J. (1990). *Human Error*. Cambridge, UK: Cambridge University Press.
- Reason, J. (1997). *Managing the Risks of Organizational Accidents*. Adershot, UK: Ashgate.
- Redaelli, I., & Carassa, A. (2015). Coordination-Artifacts Suiting: When Plans Are in the Midst of Ordering Systems. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (pp. 165–178). New York City, New York, USA: ACM.
- Reddy, M. C., Dourish, P., & Pratt, W. (2001). Coordinating heterogeneous work: information and representation in medical care. In *ECSCW 2001* (pp. 239–258). Springer.
- Richards, L., & Morse, J. M. (2012). *Readme first for a user's guide to qualitative methods*. Sage.
- Ritchie, J., & Spencer, L. (2002). Qualitative data analysis for applied policy research. In A. M. Huberman & M. B. Miles (Eds.), *The qualitative researcher's companion* (pp. 305–329). Thousand Oaks, California, USA: SAGE Publications.
- Roberts, K. H. (1990). Some characteristics of one type of high reliability organization. *Organization Science*, 1(2), 160–176.
- Rochlin, G. I. (1989). Informal organizational networking as a crisis-avoidance strategy: US naval flight operations as a case study. *Organization & Environment*, 3(2), 159–176.
- Rochlin, G. I. (1996). Reliable organizations: present research and future directions. *Journal of Contingencies and Crisis Management*, 4(2), 55–59.
- Rochlin, G. I. (2011). How to hunt a very reliable organization. *Journal of Contingencies and Crisis Management*, 19(1), 14–20.
- Rochlin, G. I., La Porte, T. R., & Roberts, K. H. (1987). The self-designing high-reliability organization: Aircraft carrier flight operations at sea. *Naval War College Review*, 40(4), 76–90.
- Roe, E., & Schulman, P. (2008). *High Reliability Management: Operating on the Edge*. Stanford, California, USA: Stanford Business Books.
- Rogers, Y. (2012). HCI theory: classical, modern, and contemporary. *Synthesis Lectures on Human-Centered Informatics*, 5(2), 1–129.
- Roland, H. E., & Moriarty, B. (1990). *System safety engineering and management*. Hoboken, New Jersey, USA: John Wiley & Sons.

- Saldaña, J. (2012). *The Coding Manual for Qualitative Researchers (2nd edition)* (Second Edition edition). Los Angeles, California, USA: SAGE Publications.
- Sandom, C. (2009). Safety assessment and human factors. In C. Sandom & R. Harvey, (Eds.), *Human Factors for Engineers* (pp. 333–347). London, UK: IET.
- Sarter, N. B., Woods, D. D., & Billings, C. E. (1997). Automation surprises. *Handbook of Human Factors and Ergonomics, 2*, 1926–1943.
- Schreier, M. (2012). *Qualitative Content Analysis in Practice*. Thousand Oaks, California, USA: SAGE Publications.
- Schulman, P. R. (1993). The analysis of high reliability organizations: a comparative framework. *New Challenges to Understanding Organizations*, 33–54.
- Schulman, P., Roe, E., Eeten, M. van, & Bruijne, M. de. (2004). High reliability and the management of critical infrastructures. *Journal of Contingencies and Crisis Management*, 12(1), 14–28.
- Schweiger, David M., William R. Sandberg, and Paula L. Rechner. "Experiential effects of dialectical inquiry, devil's advocacy and consensus approaches to strategic decision making." *Academy of Management Journal* 32.4 (1989): 745-772.
- Shorrock, S., Woldring, M., & Hughes, G. (2004). *The Human Factors Case: Guidance for Human Factors Integration* (No. HRS/HSP-003-GUI-01). EUROCONTROL.
- Shrivastava, P. (1994). Technological and organizational roots of industrial crises: Lessons from Exxon Valdez and Bhopal. *Technological Forecasting and Social Change*, 45(3), 237–253.
- Simon, H. A. (1957). *Models of man: social and rational*. Oxford, UK: Wiley.
- Sittig, D. F., & Singh, H. (2010). A New Socio-technical Model for Studying Health Information Technology in Complex Adaptive Healthcare Systems. *Quality & Safety in Health Care*, 19 (Suppl 3), i68–i74.
- SKYbrary. (2014). Safety Nets - SKYbrary Aviation Safety. Retrieved August 6 2015, from [www.skybrary.aero/index.php/Safety\\_Nets](http://www.skybrary.aero/index.php/Safety_Nets).
- Snook, S. A. (2002). *Friendly Fire: The Accidental Shootdown of U.S. Black Hawks over Northern Iraq*. Princeton University Press.
- Spradley, J. P. (1980). *Participant Observation* (1st edition). New York: Holt, Rinehart and Winston.
- Stinchcombe, A. L. (1978) *Theoretical Methods in Social History*. New York: Academic Press.
- Stolzer, A. J., Halford, C. D., & Goglia, J. J. (2008). *Safety management systems in aviation*. Ashgate Publishing, Ltd.
- Strong, D. M., & Volkoff, O. (2010). Understanding organization-enterprise system fit: a path to theorizing the information technology artifact. *MIS Quarterly*, 34(4), 731–756.
- Sujan, M.-A. (2001). *Human and organisational aspects in safety relevant technical systems*. Logos/Verlag.
- Sumwalt, R. (2014, August). *Controlled Flight Into Terrain. The Problem that Never Went Away*. Presented at the Bombardier Safety Standdown, Sao Paulo, Brazil. Retrieved June 9, 2015, from [www.nts.gov/news/speeches/rsumwalt/Documents/Sumwalt\\_140811.pdf](http://www.nts.gov/news/speeches/rsumwalt/Documents/Sumwalt_140811.pdf).
- Tamuz, M., & Harrison, M. I. (2006). Improving Patient Safety in Hospitals: Contributions of High-Reliability Theory and Normal Accident Theory, 41.
- Tarozzi, M. (2008). *Che cos'è la grounded theory*. Roma: Carocci.
- Tasca, L., L. (1990). *The Social Construction of Human Error* (PhD Thesis). State University

- of New York.
- Tiwana, A. (2004). Beyond the black box: knowledge overlaps in software outsourcing. *Software, IEEE*, 21(5), 51–58.
- Turner, B. (1976). The organizational and inter-organizational development of disasters. *Administrative Science Quarterly*, 378–397.
- Turner, B. A. (1981). Some practical aspects of qualitative data analysis: one way of organising the cognitive processes associated with the generation of grounded theory. *Quality and Quantity*, 15(3), 225–247.
- Turner, B. A. (1983). The use of grounded theory for the qualitative analysis of organizational behaviour. *Journal of Management Studies*, 20(3), 333–348.
- Turner, B., & Pidgeon, N. (1997). *Man-Made Disasters*, (2nd edition). Boston: Butterworth-Heinemann.
- Vaughan, D. (1989). Regulating Risk: Implications of the Challenger Accident\*. *Law & Policy*, 11(3), 330-349.
- Vaughan, D. (1990). Autonomy, interdependence, and social control: NASA and the space shuttle Challenger. *Administrative Science Quarterly*, 225-257.
- Vaughan, D. (1997). *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA* (1 edition). Chicago, USA: University Of Chicago Press.
- Vaughan, D. (2004). Theorizing Disaster Analogy, historical ethnography, and the Challenger accident. *Ethnography*, 5(3), 315–347.
- Vaughan, D. (2005). Organizational rituals of risk and error. In B. Hutter & M. Power (Eds.), *Organizational encounters with risk* (pp. 33–66). Cambridge University Press.
- Vaughan, D. (2009). On Slippery Slopes, Repeating Negative Patterns, and Learning from Mistake: NASA's Space Shuttle Disasters. In *Controversies in Science & Technology - Volume 2 From Climate to Chromosomes* (pp. 262–275). Mary Ann Liebert, Inc., publishers.
- Waterson, P. E., & Jenkins, D. P. (2010). Methodological considerations in using AcciMaps and the Risk Management Framework to analyse large-scale systemic failures. In *System Safety 2010, 5th IET International Conference on* (pp. 1-6). IET.
- Weick, K. E. (1989). Theory construction as disciplined imagination. *Academy of management review*, 14(4), 516-531.
- Weick, K. E. (1993). The collapse of sense making in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 628–652.
- Weick, K. E. (1998). Introductory essay—improvisation as a mindset for organizational analysis. *Organization science*, 9(5), 543-555.
- Weick, K. E. (2000). *Making Sense of the Organization*. John Wiley & Sons.
- Weick, K. E., & Roberts, K. H. (1993). Collective mind in organizations: Heedful interrelating on flight decks. *Administrative science quarterly*, 357-381.
- Weick, K. E., Sutcliffe, K. M., & Obstfeld, D. (2008). Organizing for high reliability: Processes of collective mindfulness. *Crisis management*, 3, 81-123.
- Weick, K. E., & Sutcliffe, K. M. (2011). *Managing the Unexpected: Resilient Performance in an Age of Uncertainty*. John Wiley & Sons.
- Widdowson, A., & Carr, D. (2002). *Human factors integration: Implementation in the onshore and offshore industries* (Research Report No. 001). Health & Safety Executive.
- Wiener, E. L. (1977). Controlled flight into terrain accidents: System-induced errors. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 19(2), 171–181.

- Wiener, E. L. (1989). *Human factors of advanced technology ('glass cockpit') transport aircraft* (Vol. 177528). NASA Ames Research Center.
- Willcocks, L., & Sauer, C. (2000, May). High risks and hidden costs in IT outsourcing. *Financial Times Magazine*.
- Woods, D. D. (2005). Creating foresight: Lessons for enhancing resilience from Columbia. In W. Starbuck & M. Farjoun (Eds.), *Organization at the limit: lessons from the Columbia disaster*. John Wiley & Sons.
- Woods, D. D. (2010). *Behind human error*. Ashgate Publishing, Ltd.
- Woods, D. D., Dekker, S., Cook, R. I., Johannesen, L., & Sarter, N. (2010). *Behind Human Error* (2 edition). Farnham: Ashgate Publishing, Ltd.
- Woods, D. D., & Sarter, N. B. (2000). Learning from automation surprises and going sour accidents. *Cognitive Engineering in the Aviation Domain*, 327–353.
- Yin, R. K. (2004). *The Case Study Anthology* (1 edition). Thousand Oaks, California, USA: SAGE Publications, Inc.
- Yin, R. K. (2008). *Case Study Research: Design and Methods* (4th edition). Los Angeles, Calif: SAGE Publications, Inc.



## APPENDIX A: REVIEW OF RELEVANT OS METHODOLOGIES

This appendix reviews research methodologies found in the organisational safety (OS) literature, and draws conclusions about their potential for supporting the present research. The reviewed methodologies belong to two classes:

1. **Normative methodologies.** Normative methodologies are mainly intended for safety practitioners for carrying out retrospective or perspective analysis of organisational accidents and extracting practical recommendations for safety improvement. These methodologies are based on the application of a theoretical model of organisational drift, which orients data collection and modelling.
2. **Interpretive methodologies.** Interpretive methodologies have been crafted and used by scholars with an interest in producing a theory of the organisational causes of disaster, rather than just extracting safety recommendations. As in classic sociological qualitative research, the emphasis of interpretive methodologies is on working bottom-up for developing a model that account for the data that has been collected—rather than modelling a situation in a top down fashion.

A summary of the reviewed methodologies is reported in the Table 27.

**Table 27.** Summary table of relevant OS methodologies available in the OS.

NORMATIVE METHODOLOGIES	INTERPRETIVE METHODOLOGIES
<ul style="list-style-type: none"> <li>• Management Oversight Risk Tree (MORT);</li> <li>• ACCIMAP;</li> <li>• Systemic Theoretic Process Analysis (STPA)</li> <li>• System Failure Method (SFM).</li> </ul>	<ul style="list-style-type: none"> <li>• Grounded Theory Methodology (GTM);</li> <li>• Historical Ethnography;</li> <li>• Causal Mapping;</li> <li>• Analysis of the Vocabulary of Safety;</li> <li>• HRO cases;</li> <li>• Disciplined Imagination.</li> </ul>

## **NORMATIVE APPROACHES**

### **Management Oversight Risk Tree (MORT)**

Arguably, the Management Oversight Risk Tree, or MORT, represents the earliest attempt to include managerial and organisational factors in safety analysis (Leveson, 2011, p. 31). MORT has been developed by Johnson in the mid 1970ies at the US Energy Research and Development Administration (ERDA, formerly AEC) (Kirwan & Ainsworth, 1992, p. 208). It has been developed following a review of best safety practices, found across a variety of industries, aimed at articulating general reference principles that could guide the assessment of safety management and safety assurance practices within organisations (Le Coze, 2008). In particular at the root of the analysis is a checklist that supports the systematic questioning about the state of the organisational activities intended to support the barriers and controls of hazardous work processes (Le Coze, 2008, p. 142).

A typical MORT analysis consists of reconstructing the chain of events leading to an accident, of identifying the contributing (first level) factors and the intervening (second level) factors located at the sharp end, and of identifying the systemic managerial (third level) factors located at the system blunt end (Kirwan & Ainsworth, 1992, p. 208). For each event, the analysis needs to look up in the organisation for “management and design decisions about people, equipment, processes and procedures that are relevant to the accident” (NRI, 2009, p. x). The MORT checklist supports this process by providing questions about the quality and adequacy of risk analysis processes, of maintenance plans, of specified responsibilities and task assignments, of levels of supervision, of adopted procedures and training, and so forth (see e.g., NRI, 2009).

Leveson notes that MORT is ultimately a usable method that can be applied for auditing purposes across a variety of industries. However, she also notes that “such generalisability also limits its usefulness” (Leveson, 2011, p. 31). Le Coze (2008) supports this view, and adds that MORT checklist does not allow to probe into the organisational conditions that produced a given managerial behaviour. Organisational factors were not extensively theorised at the time MORT was developed; hence they are not included in the method. Furthermore, MORT’s output consists on an explanation of how and to which extent a situation differs from best known safety practices—as embedded in the checklist—, and which of these best practices could have prevented an accident from happening (Le Coze, 2008). Certainly, this makes the method appropriate for assessing the adequacy of safety management practices found in an organisation (Kirwan & Ainsworth, 1992, p. 208), but not for interpreting and explaining a specific accident or failure situation. The method does not support in fact the reconstruction of those conceptual links between the

organisation, its history and context (Le Coze, 2008) which are of relevance in order to appreciate the processes of organisational drift into failure or adaptation.

## **ACCIMAP**

ACCIMAP is a risk management technique developed by Jens Rasmussen and Inge Svedung for the Swedish Rescue Service Agency. It aims at characterising the systemic precursors to accidents located at various levels of society (Rasmussen & Svedung, 2000). The approach is rooted on Rasmussen's socio-technical control model of risk management. Within this model, the socio-technical control system includes those actors at the blunt end that might influence the way a hazardous process is managed at the sharp end of the system. Such actors can be found at six distinct levels: government, regulators, branch association, company, management and staff (Rasmussen & Svedung, 2000, p. 68; Leveson, 2011, p. 31). Accidents arise from the unintended and undesirable combination of the side effects resulting from the stream of normal decisions of organisational actors located at these levels. Such side effects can create the landscape through which 'an accidental flow of events might evolve' (Rasmussen & Svedung, 2000, p. 24).

As in traditional accident analysis, ACCIMAP analysis starts with the reconstruction of the chain of events at the work process level (sharp end) to explain how the accident has evolved over time. However, this is just a point of departure. The analyst has to search for links between failures found at the level of the work process and the (biased) decisions taken by different actors at the various levels of the socio-technical system at different point in time prior to the accident. The last step of the analysis consists of characterising the local information environments within which these actors operated. The selection and presentation of information has in fact implications regarding 'what people will see as the problem to be solved, and what aspects of that problem are relevant and which are not' (Dekker, 2011a, p. 112). The objective is not to search for managerial errors, but for those normal decisions, taken at different points in time by organisationally 'disconnected' actors, which might have induced unsafe conditions in the work system. Ultimately, such analysis is able to report systemic sources of bias in organisational decision-making.

For instance, in the analysis of the loss of the Brazilian space programme's launch vehicle in 2003, Johnson and de Almeida (2008) identified systemic bias at government level (absence of legislation on complex system development and analysis, delays and retention in resources allocation due to resources control by Finance Ministry); company level (conflicting priorities and problems in collecting feedback about system anomalies); and staff level (absence of a risk management team).

Ultimately, ACCIMAP provides practitioners and accident investigators with a powerful unit of analysis including all of the societal levels involved in risk management. In

particular, the approach demands an appreciation of the vertical integration among socio-technical levels and its link to operational safety. This is original as usually different levels of such system are the subjects of independent disciplines.

However, it can be observed that the approach has not been developed for theorizing purposes. ACCIMAP is unable to support the analysis of those softer cultural and sense making dimensions of organisational failure—such as those embedded in the models of Turner, Vaughan and Weick (see §3.2.4)—that can explain why a group of organisational actors behaved in a certain way. For instance, the Waterson and Jenkis's ACCIMAP analysis of an infection outbreak occurred at an NHS trust in the UK between 2004 and 2006 (Waterson & Jenkis, 2010), ultimately had to import the theoretical concepts of Vaughan's normalisation of deviance to explain the managerial inability to act despite the presence of a known risk, and Weick's culture of entrapment to explain why the behaviour of clinicians and other healthcare professionals become trapped into a particular (ineffective) line of action.

### **Systemic Theoretic Process Analysis (STPA)**

The systemic-theoretic process analysis (STPA) (Leveson, 2011) is a hazard analysis methodology developed by Leveson at MIT, which is based on the system theoretic model of accident (STAMP). STAMP (§ 3.2.2) sees systems as composed of interrelated components kept in a state of dynamic equilibrium. Safety is an emergent property that is achieved when system's and components' behaviours are kept within predefined safety constraints. On the other hand, failure is explained in terms of what safety constraints have been violated and why controls have been inadequate in enforcing those constraints. The categories of control flaws include (i) inadequate enforcement of constraints (control actions); (ii) inadequate enforcement of control action; and (iii) inadequate or missing feedback.

STPA can be applied for identifying sources of hazards related to the human, hardware and software components of the operational system (sharp end). Furthermore, similarly to ACCIMAP, it can identify sources of hazards related to the socio-technical control structure (blunt end), i.e. managerial and organisational flaws.

When applied to the identification of such flaws, STPA provides to system stakeholders a structured method to identify where relevant constraints at various hierarchical levels can be lost. From here, the approach promotes the identification of the information and documentation to make sure that safety constraints are enforced in system design (Leveson, 2011). The analysis has two main steps: (i) identify the potential for inadequate control of the system that could lead to a hazardous state; and (ii) determine how each potentially hazardous control action identified in the first step could occur.

Leveson used STPA to carry out a programmatic and organisational risk analysis related to a NASA control function, the NASA Independent Technical Authority (ITA). This organisational function was created to reduce the lack of independence of the Safety Program from the Space Shuttle program manager, as this condition was found to contribute to the loss of Columbia (Leveson, 2011, p. 231). Leveson analysed the requirements that were not being enforced in the ITA organisational structure; as well as basic risks and coordination risks. (These latter usually arise when multiple actors control the same process.) Eventually, STPA was helpful to identify potential changes to the safety control structure that could eliminate or mitigate identified risks. In particular, the methodology allowed to identify about 150% more hazards than a separate expert analysis of the same function carried out by NASA.

STPA provides a structured approach to identify systemic hazards in the safety control structure and devise appropriate remedies accordingly. The identified hazards can be in fact quickly fixed by adding additional safety constraints. Also, the methodology has the merit of expanding the scope of traditional hazard analysis approaches, such as HAZOP, event tree analysis and fault tree analysis. STAMP includes types of hazards that exceed the scope of these models. When used for theoretical purposes, the approach seems to suffer from the same limitation as ACCIMAP. The approach has not been designed to account for and explain those softer organisational and social processes involved in drift into failure—such as the development of beliefs dysfunctional to safety goals. For instance, while the approach is helpful to spot missing, poorly implemented or uncoordinated safety constraints, it reduces explanations of failure to flaws in the control structure, i.e., it does not allow understanding why such flaws come into place.

*Table 28. Control flaws leading to hazard (Leveson, 2002)*

<ol style="list-style-type: none"> <li>1. Inadequate Enforcement of Constraints (control actions) <ol style="list-style-type: none"> <li>1.1. Unidentified hazards</li> <li>1.2. Inappropriate, ineffective, or missing control actions for identified hazards <ol style="list-style-type: none"> <li>1.2.1. Design of control algorithm (process) does not enforce constraints <ul style="list-style-type: none"> <li>• Flaws in creation process</li> <li>• Process change without appropriate change in control algorithm (asynchronous evolution)</li> <li>• Incorrect modification or adaptation</li> </ul> </li> <li>1.2.2. Process models inconsistent, incomplete or incorrect <ul style="list-style-type: none"> <li>• Flaws in creation process;</li> <li>• Flaws in updating process (asynchronous evolution)</li> <li>• Time Lag and measurement inaccuracies not accounted for</li> </ul> </li> </ol> </li> </ol> </li> </ol>
--

- 1.2.3. Inadequate coordination between controllers and decision makers
- 2. Inadequate execution of control actions
  - 2.1. Communication flow
  - 2.2. Inadequate actuator operation
  - 2.3. Time Lag
- 3. Inadequate or missing feedback
  - 3.1. Not provided in system design
  - 3.2. Communication flow
  - 3.3. Time lag
  - 3.4. Inadequate sensor operation (incorrect or no information provided)

### **System Failure Method (SFM)**

Developed by Fortune and Peters at Open University, the System Failure Method (SFM) draws extensively on system concepts in order to produce a systemic interpretation of accidents and disasters (Fortune & Peters, 1995). SFM belongs to the class of interpretive system approaches to management, such as Soft System Analysis (Checkland, 1999); hence it put great emphasis on the understanding of the multiple subjective viewpoints and perspectives that are involved in the development of failure. However, in comparison to other interpretive system approaches, the SFM focuses more on learning from cases of accidents and disasters, rather than managerial optimization. To Peters and Fortune, the production of case studies on failure is part of an organisational analogical learning process by which organisations can learn by comparing their own activities, processes or performances with failed safety practices occurring in their and other industries (Fortune & Peters, 1995, p. 8).

In short, SFM consists of a structured approach that allows the researcher to move from the real world failure to an abstracted representation of it. From here, system thinking, qualitative modelling, and comparison provide insights into the sources of failure, so that it is possible to derive lessons learnt and recommendations for improvement. Throughout the analysis, the SFM can accommodate a variety of qualitative system modelling techniques to capture the multiple interconnected and salient features of a situation. To Fortune and Peters, such flexibility is one of the distinctive traits over others, more rigid, investigative techniques.

In practice, an SFM analysis entails a sequence of stages, namely pre-analysis, system modelling, and comparison. During the *pre-analysis* stage, the objective is to gather and organize the information related to the situation under analysis, and produce a focused definition of failure. SFM acknowledges in fact that although a situation has been labelled

as a failure, it is necessary to devise a more specific definition as the focus of the analysis (Fortune & Peters, 1995, p. 100). For instance, the authors suggest that a major accident, such as the fire at Manchester Airport in 1999, might comprise several failures, each of which can be chosen as the exclusive object of analysis. The choice of which failure to consider for analysis depends on the purpose, viewpoints and perspectives involved. For instance, in relation to the Manchester Airport's fire, SFM's authors noted that a firm that manufactures smoke-hoods might have had an interest on commissioning a study so that its findings could be used "to support the case that the Civil Aviation Authority should require carriers to provide all passengers with smoke-hoods. For such a study the perspectives and viewpoints of the clients, passengers and carriers would be of primary importance. However, if the Civil Aviation Authority itself commissioned a system failures study of the accident with a view to examining the adequacy of existing British Civil Airworthiness Requirements, then many more viewpoints and perspectives would have to be taken into account" (Fortune & Peters, 1995, p. 93).

During the second stage, *system modelling*, the analyst will translate the insights gained from the earlier pre-analysis phase into a qualitative systemic representation. Within SFM, failure is in fact conceived as an "output, or lack of outputs, of transformations processes carried out by a system" (Fortune & Peters, 1995, p. 101). Consequently the understanding of failure requires the representation of the broader systems within which such transformations have occurred.

To note that both the pre-analysis and the system modelling phases can benefit from the use of a variety of qualitative systemic modelling techniques, such as a (i) spray maps, (ii) rich picture, (iii) relationship diagrams, and (iv) multiple cause diagrams, for the pre-analysis, and (i) input-output diagrams, (ii) system maps, and (iii) influence diagrams, for the analysis. According to the authors, such flexibility allows the collection of depth insights into the interconnected features of a situation.

The subsequent stage, *comparison*, consists of comparing the model of the actual situation with existing "models of how a situation should be operated and managed if it is to be operated without failure" (Fortune & Peters, 1995, p. 110). This phase is at the core of the model since here insights about the potential sources of failure are identified. The first comparison is carried out at system level against an idealized model of a robust system. Such representation, called the Formal System Model (FSM), embeds different components, interactions and system boundaries that are supposed to denote the conditions under which a system produces the intended output. By comparing this idealized model with the model of the actual situation under analysis it is possible to identify discrepancies in the latter. Such discrepancies might pertain specific parts, and can be investigated through subsequent focused comparisons with models of good

practices, related to aspects of communication, control, and human factors at the individual, group and organisational levels. For instance, concerning the analysis of organisational factors, the SFM suggests the model of planning of Hall (1982) as a basis for comparison. This model explains failures based on interactions among a triangle of actors, namely bureaucrats, politicians, and the public.

IN addition to this model, SFM comes with an array of theoretical models available in the literature which can be used in the comparative stage of the process. Such models are those that have been found most useful through several applications of the model. However, this base is only suggested, and the authors maintain that future uses of the SFM can include more models.

## **INTERPRETIVE APPROACHES**

### **Grounded Theory Methodology**

The Grounded Theory Methodology (GTM) (Glaser & Strauss, 2009) lies at the core of the seminal work of Barry Turner, who is usually referred to as the first scientist who has adopted an organisational perspective on the study of accidents (Catino, 2013; Dekker, 2012). His research on safety-critical systems—initiated in the mid-seventies—has been arguably the main theoretical precursor to the contemporary safety debate on organisational safety that developed since the early 90s (Catino, 2002). Up to then, if we omit the Perrow's NAT and the High Reliability Organisations (HRO), accidents were conceptualised mainly as a sequence of technical failures and human errors— but the societal components of disaster were left mainly unexplored.

Turner's Man Made Disaster theory has resulted from a GTM analysis of eighty-four accident reports published by the UK government over the period 1965–1975. The reports were compared in order to extract common underlying themes across different accidents (Fortune & Peters, 1995, p. 48). To Turner, relevant themes consisted of organisational behaviours that hampered the organisation ability to control risk. These included for instance poor handling of critical information, incorrect assumptions, lack of compliance to regulations, reluctance to accept danger. These themes emerged after an initial analytical step centred on the comparative analysis of three accident reports only. Subsequently, Turner explored if and how the identified themes applied to the other accidents.

By following this approach, Turner was able to produce a six-stage longitudinal model of accident disaster development, where the protracted incubation period was the most important phase. In fact, by characterising this process, he has been able to show for the



first time that accidents are preceded by a history of anomalies and hazards being downplayed.

Turner wrote extensively on the GTM data handling procedures he used (Turner, 1981, 1983). GTM is a structured approach to social inquiry to generating theories of social interaction from qualitative data. It has emerged in the 60ies from the seminal work of Glaser & Strauss (2009) as an alternative to the mainstream quantitative positivistic sociology, which emphasised hypothetic deductive testing of existing grand theories of society, while relegating the role of qualitative and explorative research to the margins of social research.

The development of GTM proceeds by developing systematic conceptualisations: the researcher seeks to construct a theory by extracting concepts from incidents and events found across the data, and establishing conceptual connections between them until a single theoretical explanation is achieved (Tarozzi, 2008; Richards & Morse, 2007; Cicognani, 2002). Operationally, this happens through a set of systematic steps through which data is dissected, conceptualised and reassembled in new explanatory ways (Flick, 2009). Initially, different data segments are coded for individual conceptual categories which are very close to the data, and do not clearly relate nor to each other, nor to any existing theoretical framework—under GTM the researcher seeks to avoid imposing existing theories or ideas. As the research progresses, while some of the generated categories have to be dropped, links between other categories may become clearer, until the data can be reduced to a few conceptual categories—more abstract and general compared to those generated in the earlier phase—that provide a theoretical explanation of what the data stands for.

Two strategic trademarks of this analytic process are the *constant comparative method* and the *theoretical sampling* (Partington, 2002). The constant comparative method consists of comparing continuously along the research process events, properties and dimensions found across incidents. For instance the researcher can compare evidences about incidents with evidences from other incidents previously collected and grouped under a concept that stand for those data. The comparison will generate insights about whether the new incident can be allocated under the existing concept, about whether the concept has to be refined in order to account for the new incident, or if a new concept will have to be provided. Such a comparison is fundamental for the analysis as it allows the researcher to differentiate between categories and find dimensions and properties peculiar to the specific category.

Theoretical sampling is the process by which the researcher maintains control over the emergent theory. Through theoretical sampling it is possible to plan the next phases of

data collection in order to minimise or maximise differences between cases. This approach differs from statistical sampling in that the focus is not on generalisation (i.e., collecting cases which abide to the law of statistical inference), but on extending or refining the developing theory. An aspect of theoretical sampling directly borrowed from Bacon's inductive method is the exploration of cases found in the data that seems to contradict the emergent theory. The analysis of negative cases is intended to strengthen the final emergent theory.

### **Historical Ethnography**

The work of the American sociologist Diane Vaughan brought interpretive sociological analysis to the realm of safety. Her analysis of the Challenger and Columbia disasters departed from Turner's observation that accidents are usually preceded by history of signals about anomalies being downplayed. However, with respect to Turner, Vaughan has drawn on extensive *micro* data about key decisions preceding launch as well as *macro* data about institutional and organisational dynamics. This is original since Turner did not have macro level data to explain decision-making (Vaughan, 1997), but mainly accident reports.

Also, Vaughan's work has important precursors in the work of Star and Gerson (1987), who studied how unexpected anomalies are responded to in scientific work. These scholars have found that anomalies' definition, negotiation and control are influenced by the institutional context and by the controls that have been put in place to deal with them. To Star and Gerson, the response to deviant or discrepant events and the consequential "trajectory" of these anomalies does not happen in isolation but is a reflection of the organisational system within which it is embedded.

Vaughan approached her investigation of the NASA disasters of Challenger and Columbia historically. She departed from the accident to reconstructs the "trajectories" of the O-ring and the debris tiles anomalies within NASA prior to the disasters, and from the rationale of the insiders. In fact, anomalies prior a disaster might or might not be interpreted as a safety threat within the organisation, might be clear or ambiguous, might be perceived as imminent danger or not (Edmondson et al., 2005).

The American sociologist concluded that prior to the event, personnel operated according to the rationale embedded in the document and NASA organisational culture, thus the disaster was not a case of negligence as firstly stated in the official presidential report; rather it was a case of organisational-systemic failures—namely the failure to learn, the failure to aggregate relevant information, the regulatory failure to overseeing safety-critical processes (Vaughan, 1989, 1990, 1996). Eventually, she has been able to explain accident causation as consequential to three organisational patterns—normalisation of

deviance, structural secrecy, and conflicting institutional objectives—which systematically can lead organisations to downplay risk while pursuing economic targets (Vaughan, 2005). Ultimately, the explanations generated by her approach are structural in that it explains actions of the past as constructed, formed, or organized by major institutional forces, rather than unconstrained individual choices (Parsons, 2007).

She qualified her approach as an historical ethnography, i.e., a way to elicit structure and culture from documents created prior to an event in order to understand how people in another time and place made sense of things” (Vaughan, 2004). Vaughan based her conclusions on the collection of a wide amount of interviews with NASA engineers, contractors and managers and official documentary data over a period of nearly a decade. This material was scrutinised in order to understand how individual sense making, cultural understandings, and actions prior the accident were shaped by historical institutional, ideological, economic, political, and organisational forces (Vaughan, 2004).

Also, Vaughan provided a discussion about the process of theorizing from empirical data. She states that researchers should theorise by comparing analogous events, activities, or incidents found in different social settings, of various complexities, sizes, and functions. She legitimates the use of this process, that she calls *analogical theorising*, by observing that it is commonly found in many classic sociological works (she mentions Blau, 1964; Goffman, 1961; Hirschamn, 1970; Hughes, 1984). To her, only Stinchcombe (1978) refers openly the search for analogies and differences across cases, but for similar social units only (e.g., all nation states). But a second and more important aspect in analogical theorising is perhaps the idea of linking patterns found in the data to known theories or concepts as the research progresses. The rationale being that there are always concepts, models, and theories influencing researcher interpretations. Making them explicit makes possible to reject, conceptualise, and/or working toward more generable explanations.

### **Causal Mapping**

In the book *Friendly Fire*, Snook (2002) produced a revisionist account of the accidental shutdown of two U.S. Black Hawk helicopters by a U.S F-15 fighter jet in Northern Iraq in 1999. In the same book, the American author provides a succinct three pages description of his analytical strategy.

Snook could draw on a large evidentiary base including both primary and secondary data. Time-stamped audio, video and data tapes were coupled with in-depth interviews with multiple subjects located at various positions in the army ranks. These interviews helped to understand “what information was available to whom and when in the build up to the accident.”(Snook, 2002, p. 18)

Drawing on this data, he worked through a five steps process in order to construct a causal map of the tragedy. Step one consisted of constructing a precise multileveled timeline of the events proximal to the tragedy. The timeline showed interactions across the crews involved (the two Black Hawks, the two F-15, and an AWACS) prior to the shutdown. Step two coupled this map was with the practical accounts emerged by the in-depth interviews he collected from multiple informants. So, while step one provided a pure factual description of the events, e.g., “At approximately 1030, the lead pilot fired an AMRAAM missile at the trail helicopter”, step two added various institutional and more general explanations, e.g., “The accident resulted from a failure to integrate helicopter operations into OPC flight operations.”(Snook, 2002, p. 18). Step three consisted of systematically subjecting this early version of the causal map to a series of questions, in order to add details, sharpen temporal and causal dependencies between the events, and developing a theoretically oriented explanation. During this step, Snook reported to have drawn on “a rich set of data, a working knowledge of military organisations, and broad reading of behavioural science literature” (Snook, 2002, p. 20). Also, he has made references to the methodological concepts and procedures of ‘disciplined imagination’, ‘theoretical sensitivity’, ‘open’ and ‘axial’ coding, mentioned by other OS scholars; although a clear description about these and the way they were used in his research is not provided. In Step four, the different parts of the map were subsequently subjected to counterfactual (“what if”) interrogation, in order to challenge the causal significance of each major event. This process allowed to assess whether the events involved were actually significant to the development of the history as it happened, or not (in this latter case they were removed from the map). The final steps consisted of developing a single story line that could stand for the events and causal links embedded in the map, and validating the final story against the data.

Overall, Snook’s methodological account appears to be very pragmatic. Further, his constant concerns with systematic questioning and counterfactual analysis seem to strengthen the confidence on his results. It is however difficult to understand fully how he abstracted the first description of events (Step 1) into a broader multileveled conceptual explanation (Step 3). Although he made references to various methodological concepts, such references were only indicative: arguably more information would have been beneficial to understand how such concepts have been implemented. Such a lack of information can be understood considering that his methodological account was part of a book intended for a broader audience including practitioners, and not just academics.

## **Analysis of Vocabulary of Safety**

William Ocasio, a Professor in management and organisation at the Kellogg University, examined the role of safety culture in the production of the Columbia Accident (Ocasio, 2005). His work departed from CAIB key conclusion that safety culture was a key factor in the development of that disaster. Through a in depth linguistic analysis, he came to confute the CAIB conclusion.

Ocasio's methodological approach builds on a theoretical view of language in organisation according to which specialized domain of organisational practices are articulated by systems of specific and interrelated vocabularies (Ocasio, 2005; Loewenstein & Ocasio, 2004, p. 4). For instance, within organisations it is possible to find vocabularies of business strategy (e.g., strategic planning, strengths-weaknesses-opportunities, sustainable competitive advantage, core competence, experience curve), of corporate human resources (e.g., workout, vitality curve, Six Sigma), and safety and risk (e.g., disaster, risk management, acceptable risk, system anomaly). Each of such vocabularies is composed by a set of interrelated words or group of words for articulating specific areas of organisational practice (Loewenstein & Ocasio, 2004, p. 5). Vocabulary of organising might include references to "formal structure, to organisational roles, processes, systems and techniques employed in the formulation and implementation of organisational practices" (Loewenstein & Ocasio, 2004, p. 5).

An important precursor to this view of language is the classic work of March and Simon (1958), according to which organisation's vocabularies and classification schemes embeds the concepts that frame the way a given issue will be interpreted and discussed. Issues that fit within such system of concept are communicated promptly within the organisation; issues or problems that do not fit, are communicated with difficulty. The system of concept reflected in the organisation vocabulary influences the way in which organisational members perceive and interpret the world.

One implication of these considerations is that while verbal expressions and their linkages are indicative of the culture of an organisation, the examination of vocabularies 'provides an opening to understand what organisations think (Douglas, 1986, p. 102). For instance, an analysts can trace the meaning of specific works of a given vocabulary, to observe how this varies over time and across different organisational members.

In his analysis of the Columbia accident, Ocasio used archival and historical sources to reconstruct as accurately as possible the phenomenological experience of organisational members in the NASA space shuttle program (Ocasio, 2005). By phenomenological experience, Ocasio appears to refer to the viewpoints, perceptions and interpretations held by NASA personnel prior to the tragedy. Ocasio analysis proceeded exploratively by

retrieving all documents containing key terms such as safety, risk, disaster, acceptable risk, safety of flight (Ocasio, 2005, p. 106 ). He then explored the meaning attached to these vocabularies, if and how they evolved over time, how meanings influence the categorization of anomalies, if and how existing organisation schemes were able to accommodate for uncertainty and ambiguity.

As anticipated earlier, Ocasio's findings contradict the findings of CAIB on one important point. To Ocasio, Columbia did not have a breach in safety culture, i.e., normalisation of deviance, as a main root cause. He found that although the foam debris anomaly was a known issue within NASA prior to the disaster, it was not anticipated that such event could damage the ultra-resistant carbon panels located in the left wing leading edge of Columbia. None of the risk and safety documents he analysed identified the impact of the foam debris on the reinforced carbon panels as a relevant safety issue. From his analysis, it appeared that the underlying assumption was that although the foam debris could have damaged the thermal tiles covering the Shuttle structure, it would not have breached the underlying carbon panels; and this explains why NASA considered the risk of foam debris as an 'acceptable risk'. Thus, Ocasio concluded it was the failure to establish a conceptual link between the foam debris and their impact the carbon panel as the cause of the accident. In his view, this was a case of failure of imagination, but certainly not a case of normalisation of deviance (§ 3.2.3), as the risk in question was not known.

Thus, Ocasio concluded that it is the loose coupling between two conditions the cause of the disaster. He noted that this idea seems to contradict Perrow's NAT theory (see § 3.2.5), which maintains that the potential for accident is created by high interactive complexity and tight coupling. However, the contradiction is only apparent because NAT refers to the coupling existing in the actual software, hardware and human components of the operational system; while the coupling eloquently documented by Ocasio refers to the conceptual view over coupling between two events as maintained by NASA participants.

Ocasio's methodology has the merit of bringing organisational cultural analysis to the realm of organisational safety analysis. His work seems to develop along the line of Vaughan historical ethnography (described earlier in this Appendix). Both appear to maintain purposefully an historical focus over the evolution of anomalies in organisations. However, compared to Vaughan, Ocasio embeds a more articulated epistemic foundation of language, that ultimately provide the researcher with a more empirically tractable unit of analysis, i.e., specific words and expressions as embedded in a given vocabulary of safety.

## HRO Cases

While the methodologies reviewed in the previous sections deal with after-the-fact investigation of failure, this section reports on the work of the High Reliability Organisation (HRO) scholars, which notably have investigated how high-risk organisations maintain safety (§ 3.3). HRO scholars have maintained an exclusive interest on understanding those organisational traits—such as structural patterns, type of management and operational practices, organisational culture—involved in ensuring and maintaining error free operations in safety critical domains. Their work was motivated by the absence in the literature of a theory which could explain how some kind of high risk organisation can actually operate nearly error free. Consequently, their focus was on the understanding of normal day-to-day operations in an attempt “to draw lessons from cases where operations have gone mostly according to plan, in contrast to studies that review past accidents to determine what went wrong” (Bourrier, 2002).

In particular, HRO scholar’s empirical work focused on studying such normal operations within three organisations: the US Federal Aviation Administration’s Air Traffic Control System; Electric Operations and Power Generations Departments; the peacetime flight operations of two U.S. Navy aircraft carriers. These organisations were selected as case studies given their impressive safety records. They were analysed through in-depth case studies, each case centred on a single organisation and aimed at getting insights into the specific organisational processes and strategies at play.

Despite their influence on the safety literature, it has been noted that HRO scholars have somehow underreported their methodological approaches (Bourrier, 2011). In comparison to most of the work reviewed in this chapter, they did not seem concerned with providing in-depth accounts about their methods in dedicated publications. Their rare methodological accounts cover only a limited part of the methodology.

Rochlin has provided a short general overview of the methodological orientation of her and his HRO’s colleagues, reporting that their research program “evolved from straightforward interview and survey work to a more complex blend of organisational analysis, studies of organisational culture and ethnographic observations at all levels of the organisations” (Rochlin, 1996, p. 55). More recently, he has added that HRO field-work was “as intimate as that characterising participant observations”, although participation was not possible due to the nature of the risk involved in the observed practice (Rochlin, 2011). It transpires from these observations a methodological concern with depth of the inquiry, as, arguably, the viewpoints and perspectives of practitioners were deemed to be relevant by HRO to understand how a particular organisation achieved safety.

Roberts (1990), in her investigation of peacetime aircraft carrier operations, reported on

the logistic aspects of data collection. She reported that researcher collected data by going to sea “intermittently for periods of five to ten days, making observations and learning jobs done aboard the ships” (Roberts, 1990, p. 164). During these periods, researchers entered field notes “into computers every few hours when the pace of the ship’s activity permitted”, and different researchers had a different background and were assigned to different observation location in the ship to avoid individual research bias (Roberts, 1990, p. 164). So, implicit in this note, is a HRO methodological concern with maintaining objectivity by ensuring the cross-checking of data among researchers.

Roe and Schulman (2008) have indulged into one of the longest and most detailed methodological description found in the HRO literature. In their monograph, *High Reliability Management: Operating on the Edge*, which investigated the intricate operations of California Independent System Operator (CAISO), they provided a two pages description of their methodology. First, they approached key CAISO officials in order to get permission to interview key participants. These were selected using the *snowballing technique*, in which new key participants for interviews were identified based on suggestions from previous interviewees. This went on until the research reached a point at which new participants mentioned very similar sort of issues as previously stated by previous participants (Roe & Schulman, 2008, p. 225). In general, interviews were reported to last about one hour or more, and were supported by a structured questionnaire allowing for open-ended answers. Ultimately, Roe and Schulman reported to interview about a hundred operators, and have spent an equal number of hours doing fields observations of control room operations. Unfortunately, not much is said about the qualitative data analysis techniques deployed in their study.

Overall, these rare methodological accounts provided by HRO authors appear to cover mainly the type of data and the data collection approaches used in their case studies. In particular they seem to have prioritized the collection of ethnographic notes and observations, and individual and group interviews with participants in order to maximise insights into the specific processes under study. Because of this, Bourrier has observed that HRO cases are essentially ethnographic in nature, and also that they have the merit of having brought the ethnographic approach to the study of high-risk organisations (Bourrier, 2011, p. 14).

Perrow has objected that the excessive proximity to the organisational practice under study intrinsic in the HRO research might have compromised the objectivity of HRO findings. However, this observation seems to ignore that secondary data reporting on how safety is achieved during normal operations was not available within the organisations HRO investigated. Usually, such a reporting activity is not a standard practice as it is the production of official reports of major accidents and disasters—the kind of secondary data



used by Perrow and other scholars. Consequently, HRO scholars had no choice but to gain access and collect primary data about practitioners' activities and viewpoints in order to investigate safety enhancing organisational processes. Furthermore, to further mitigate Perrow's criticism, they strived for objectivity by ensuring constant cross checking of findings.

Note however that HRO scholars seem to have remained silent about data analysis. Not account or reference has been found about the theorizing process. A notable exception to this trend is provided by Karl Weick, an eminent organisational scientist who has often been associated to the HRO group, and who has reported extensively on the theory building process in the context of organisation research. His approach, *disciplined imagination*, is discussed next.

### **Disciplined Imagination**

Weick's methodological contribution departs from the observation that theory construction has often been described as a linear problem solving process (Weick, 1989). However, this view misses the essence of the thinking involved: theory building is characterised by simultaneous parallel thinking rather than sequential, and most importantly it has at its core the imaginative exploration of metaphors and images (Weick, 1989).

To Weick, theory building is essentially a sense-making process (Weick, 2005). The need for a theory arises from a gap between available concepts and the observed reality. The scientist is called to make sense of a reality that cannot be explained using available theories. Such a conceptual gap can be filled by relying on systematic analogical reasoning: i.e., understanding new uncertain situations by projecting meaning from analogous more familiar ones. While such an analogical and sense-making process is usually intuitive and unconscious, Weick made it explicit in order to present it as a legitimate and structured approach to theory building.

Such approach, called *disciplined imagination* (Weick, 1989), relies on the design and conduct of imaginary experiments, where a researcher, confronted with its data, tentatively explores the plausibility of different metaphors in representing and expressing the complex organisational phenomena embedded in that data. Such imaginary experiments stimulate the production of various metaphorical images, which in turn are selected through careful judgment, and possibly retained for further theorizing.

For example, a notable Weick's writing elaborated a view of organisational improvisation as jazz (Weick, 1998). Here, organisational improvisation is viewed as "performative in nature, guided by technical structures and minimal social structures and involving

simultaneous reflection and action, simultaneous rule creation and following, continuous mixing of the expected with the novel” (Cornelissen, 2006, p. 15). In order to get to such a unifying perspective over organisational improvisation, Cornelissen has suggested that several tough trials were carried out by Weick, as a result of which many alternative metaphors were discarded (2006, p. 35).

The same analytical process was behind the generation of a theoretical view of organisational behaviour as a collective mind. Weick exploited this metaphor to produce an organisational explanation of an accident occurred in 1949, in which thirteen smokejumpers perished during the firefighting operations in the valley of the Mann Gulch, Montana (Weick, 1993). While early explanations attributed the disaster to the presence of extraordinary environmental factors, Weick explanation identified the collapse of sense making in the firefighting team as the sociological cause of the disaster. It was the dissolution of the system of roles, a loose leadership, inefficient communications, and biased assessment of the fire dimension that ultimately created a desegregated organisation system unable to make collective and shared sense of an emergency situation (Catino, 2002, p. 86).

Both Weick’s analyses have exploited an underlying metaphor to produce vocabularies and theoretical representations that have provided a fresh and insightful perspective over the organisational phenomena under analysis. Furthermore, such vocabularies and languages have also been referred to, debated and tested by other works (Cornelissen, 2006, p. 17). For instance, Weick himself exploited the metaphor of the organisational processes as a collective mind to explain the collective behaviour of pilots during flight deck operations (Weick & Roberts, 1993), and how HRO organisations organize around failure in a way that promotes a an ongoing state of mindfulness (Weick et al., 2008).

It must be noted that in the organisational and management science another notable scientist has raised the explicit and systematic use of metaphor to the status of to that of a legitimate research method. Similarly to Weick, Garreth Morgan’s seminal books *Images of Organisations* (Morgan, 2006) discussed the use of metaphors as devices for the interpretive reading of complex organisational realities. He proposed that organisations could be seen as machines, as organisms, as brains, as cultures, and as political systems. The choice of which metaphor to use to make sense of a given situation influences the nature of the meaning that will be projected on it. Morgan’s book enhances Weick’ disciplined imagination in two important ways. First, it suggests that the reading of a situation might require more than the selection of a single overarching metaphor: the insights of other supporting metaphors might be considered to enhance the analytical process and the ensuing explanation.

Also, Morgan appears to be more concerned about providing guidance about the overall research process. To Morgan, as in classic ethnographic research, a researcher needs first to get inside the organisation to understand a situation as far as possible on its own terms. He or she should adopt the role of the learner rather than that of an expert, in order to purposefully suspend judgment, avoid premature closure and leave room for new insights to arise. Along the process the researcher has to document the (i) events in question, (ii) what is being said by study participants about these events, and (iii) the researcher own interpretations about the 'reading' of the situation. These measures are helpful to avoiding the risk of premature closure.

## **SUMMARY**

This appendix has reviewed investigative methodologies useful for the investigation of organisational drift into failure. The reviewed methodologies belong to two classes: normative and descriptive methodologies. Normative methodologies provide a structured top down approach to the identification of safety-criticalities that did not work in a given situation. Methodologies belonging to this class mandate that data should fit the model.

Normative approaches are appropriate when the priority is on extracting recommendations for safety improvement. This is the case of safety practitioners which are moved by very practical aims as they are interested in gaining useful knowledge in the shape of a model which can generate understanding into (organisational) sources of risks. In these cases the focus is more on what safety remedies can be obtained for a given situation. Ultimately, these situations value primarily safety practical achievements above theoretical developments. As in classic engineering science, the relevant question behind normal methodologies is "Does this specific model of organisational risk work?". More specifically, "How good is this particular model at providing effective countermeasures?" when applied to the analysis of a particular situation.

Also interpretive methodologies have served the purpose of increasing our ability to understand and control risk. However, compared to normative methods, these methodologies seems to commit to a view according to which the production of knowledge is a legitimate goal in its own terms. In our review, this reflexive orientation has been highlighted explicitly by La Porte, an eminent HRO scholar, by Bourrier in her review of HRO scholars, and by Laporte, and Vaughan, who has suggested that organisations should open to ongoing ethnographic analysis giving regular feedback.

It is possible to note that the majority of the reviewed approaches, both normative and

formative, can be considered as belonging to the class of case study research. Case study research is usually defined as a research strategy allowing a researcher to concentrate on a specific situation, bounded in time and space, which is selected for intense investigation in its naturalistic settings because it contains some elements worth discovering. It aims at providing a detailed multidimensional picture of the situation that is being studied. To Banbasat: “Case study examines a phenomenon in its natural setting, employing multiple methods of data collection to gather information from one or a few entities (people, groups or organisations). The boundaries of the phenomenon are not clearly evident at the outset of the research and no experimental control or manipulation is used” (Banbasat, 1987, p.370). It can be noted that all of the defining traits of case study research, in depth examination of a phenomenon in its naturalistic context, absence of experimental control conditions, flexibility in data collection, equally apply to the investigative methods reviewed in this section—both normative and descriptive.

It can be noted however that most of the reviewed methodologies can be classified as Failure Cases, as they have been used for retrospective or perspective investigations of organisational accidents, HRO can be referred to as Normal Operation Cases, as they have focused on how organisations achieve error free operations during everyday operations. In the cases of Normal Operation Cases there is no need for the researcher to separate between the rational held before and after the event.

Noteworthy, an exception to the use of case study research is offered by Turner’s comparative analysis of 84 industrial accidents occurred in the UK. Turner analysis was based on the GTM, and at its core had the systematic comparison of official accident reports with the purpose of extracting common themes among them, which in turn became the building blocks of his ensuing Man Made Disaster Theory. GTM is well versed for such kind of analysis: it provides a structured and systematic approach to comparative analysis of large samples, such as that of Turner, supporting the progressive and controlled reduction of data to mid-level theories able to stand for that data. In the case of Turner, GTM provided a very reliable theory, considering the success that its theory has enjoyed for many years.

GTM is very data demanding, it requires cases of incidents that can be compared. In the investigation of drift, each case of incident is represented by the specific case of failure in a particular organisation. Hence the production of an incident requires an investigation first. GTM, applied to study of drift, requires the availability of cases of incidents.

It can be noted that the reviewed investigative approaches have focused on the understanding of disasters and accidents and not HAI issues. This was expected, considering that the phenomenon of drift has been conceptualised in relation to disasters

and accidents, and not to other types of failure. The only exception to this trend is represented by the Fortune and Peter's SFM method, which has been used to investigate information system failure. The method however, belonging to the class of pragmatic approaches, does not seem appropriate for the exploratory purposes of this study.

## APPENDIX B: PHILOSOPHICAL FOUNDATIONS

---

It has been acknowledged that safety science in general, and not just organisational safety (OS), lacks an explicit reflection about its philosophical foundations (Le Coze, Reiman, & Pettersen, 2012). Such a gap has been attributed to the relative novelty of the area, which does not enjoy a long tradition compared to more classic disciplines such as physics and sociology. Furthermore, its practice-oriented nature tends to favour theoretical debates about models of failure and related remedies rather than about the ontological, epistemological foundations upon which to base research and the development of these models and remedies.

Thus, in order to define the philosophical orientation of the present thesis it is useful to consider the classic division between positivistic and interpretivist research which applies to general organisational, management and sociological research. This section will present these two approaches and then will explain why an interpretive approach appears to provide a better context for the investigation of the organisational precursors to HAI issues. Positivism and interpretivism will be necessarily presented in a succinct form, as it is impossible to do justice in a few pages to the long tradition that both positions enjoy in the history of science.

### POSITIVISM

---

Positivism is the paradigm most often associated to the scientific research in Western society. Burrell and Morgan define it as an epistemology that attempts to explaining and predicting “what happens in the social world by searching for regularities and causal relationships between its constituent elements” (Burrell & Morgan, 1979). The purpose is to produce knowledge in the form of laws like statements about regularities and causal relationships (Hirschheim, 1985, p. 83). Such laws are regarded as general, in that they apply to a wide range of situations, and universal in that they are valid across time and space (Blaikie, 2007, p. 111).

Positivism embeds a strong realist and mechanistic view of the universe. It assumes the existence of an external universe existing independently of human consciousness. Such a universe is composed of distinct elements, whose behaviour can be explained through the laws of movement (Capra, 1997). The purpose of the scientist is to discover these laws, and this is possible through (i) *reductionism*, i.e., by considering that the aggregated behaviour of the whole could be reduced to those of the constituent parts, and (ii) *isolationism*, i.e., by considering that the behaviours of these constituent parts could be studied independently of each other’s and of the whole (Capra, 1997; Biggiero, 2011, p. 8). These assumptions are also at the root of classic approaches to safety and have been

fiercely criticized by safety theorists as they inhibit opportunities for safety learning (e.g., Dekker, 2011; Hollnagel, 2014).

*Objectivism* is another important positivistic assumption. According to it, objects have meaning in themselves. In other words, objectivism sees meaning and truth as immediately accessible through observation. Such a position has originated in classic Greek philosophy, has been absorbed by Scholastic realism throughout the Middle Ages, and has achieved its maximum expression in the age of Enlightenment (Crotty, 1998, p. 42). Applied to OS, objectivism suggests that direct and objective observation of relevant organisational and social facts is possible.

Viewed from a positivistic stance, the research process consists of positing a set of beliefs about regularities or causal relationships and then subjecting it to empirical testing (Hirschheim, 1985, p. 83). Implicit in this view is Popper's logic of science, based on the systematic proposition of new theories, tentatively formulated as solutions to research problems, which are then subjected to rigorous and systematic experimental testing before being accepted as valid (Blaikie, 2007). So, the experiment is regarded as the most rigorous research method, given its emphasis on the controlled manipulation of the relevant variables, and the separation between the researcher and the studied people (Corbetta, 2003).

Ultimately, positivism renders a linear and quantitative image of science, bolstered by the tenets of operationalization of variables, their quantification and generalization (Corbetta, 2003). As noted by Hirschheim, this positivistic view of science permeates so intensely our society "that knowledge claims not grounded in positivistic thought are simply dismissed as unscientific and therefore invalid" (Hirschheim, 1985, p. 83).

## **INTERPRETIVISM**

---

Interpretivism is a research paradigm having at its core the idea that the understanding of reality requires interpretation, not just observation as for positivism (Corbetta, 2003). From this perspective, social science cannot be studied in a similar manner to the natural sciences. In particular, the study of social reality requires the consideration of interpretive processes, i.e., sense making and meaning construction, since the social world that can be known is the world of the meanings attributed by organisational members that live that world (Blaikie, 2007). Such interpretations and perspectives vary depending on the interpretive and cultural frames held by each particular observer. Thus, interpretivism does not assume a single absolute social reality equal to all people.

Interpretivism has been developed in relation to sociology and organisation theory for at least three reasons. First, sociological and organisational phenomena entail a complex

of many factors that may simultaneously apply, overlap, and cannot be directly observed. Second, the selection of relevant variables is a much more subjective assessment than in the case of the physical sciences (Dennis, 2005). Third, and most importantly, it is human purpose that determines courses of action. Organisations are in fact intentional systems, i.e., system whose behaviour is governed by human motivations and purpose rather than by the laws of physics, as it is for instance in the case of engineered technical systems (e.g., Naikar, 2013; Checkland, 1999). So, although behavioural patterns might be empirically documented, understanding and prediction are weak without first understanding human motivation and values related to those involved in the situation under study.

Viewed from an interpretivist perspective, the research process requires the researcher to get immersed first within the social reality under study, for empathizing with study participants and understanding their motivations from the *inside out*—and not from the *outside in* as positivist researchers would do. The development of a theory requires the understanding “of the social world that people have constructed and live in” (Blaikie, 2007, p. 124) and not the imposition of categories imported from the outside. These are inevitably only poorly representative of the specific social reality under investigation.

Because of this focus on knowing the social reality in depth, ethnography, grounded theory, and case study are the favoured research methodologies for interpretive work in the social sciences. These methodologies promote the development of emerging theories based on an in depth understanding of “actors’ language, meanings, and accounts in the context of everyday activities” (Blaikie, 2007, p. 89). The resulting, emerging theory consists of an interpretive account provided by the researcher, which aim at developing *ideal types*, i.e., abstracted models that have been constructed by the researcher to make sense of the particular complex social reality.

Interpretive research is much more cautious about the status of research outcomes than positivism. It acknowledges that the regularities contained by the resulting model do not enjoy the status of universal, generalisable laws as for positivism (Corbetta, 2003) . Rather they consist of statement of likelihood in the form of, if A occurs, then most of the time also B will occur. Furthermore they might be restricted to the specific case under investigation.

## **RATIONALE FOR THE CHOSEN PHILOSOPHICAL POSITION**

---

In choosing between positivism and interpretivism as the philosophical foundation better suited for the present research, the following two considerations apply:



1. First, positivist research would become manifest by means of a classic quantitative method, such as controlled laboratory experiment, survey, or simulation. However, none of these methods seems feasible in the present case. These methods require in fact the availability of variables that are clearly defined and measurable. This is clearly not the case in the present research. Furthermore, the “gold standard” quantitative method, the controlled laboratory experiment, would encounter serious issues of feasibility. Laboratory experiments consist of studying intensely a small set of variables in laboratory simulated settings in order to be able to draw generalisable statement applicable to real life situations. While appropriate when investigating individual or team behaviour, experiments become more problematic as the unit of observation—individual, team, department, organisation, etc.—grows. Specifically to investigating organisational drift into failure, it is not clear how normal organisation and administrative behaviour could be reproduced in a controlled laboratory environment without resorting to extreme oversimplification and isolation from most of the variables that exist in real life. Experiments require in fact to deliberately divorce the phenomenon of interest from its context, in order to study a few variables in a controlled environment (Yin, 2009, p. 18).
2. Second, interpretivism seems particularly appropriate as it considers the insiders viewpoints and perspectives of organisational members, i.e., important precursors to organisational failure and success. As discussed earlier, objectivism is one of the main epistemological assumption associated with positivism. Objectivism asserts that meaning and truth are immediately accessible by means of observations. This assumption seems particularly problematic for the present research as it does not consider the multiple interpretations that guide action and decision making in organisations. As discussed in section 3.2, the precursors to organisational failure and successful safety performances include the (multiple) organisational perspectives and frames and viewpoints that develop inside the organisation. This calls for a research paradigm that considers intensely such views, as reflected in organisational members’ intentions, language and meaning. As discussed earlier, this is clearly what interpretivism does. Applied to OS, this paradigm has at its core the comprehension of how organisational members experience and explain the safety-critical systems they are operating or managing (Pettersen, 2008, p. 94). In interpretive research, the researcher has to wear the learner hat in fact, so to sympathize with the interpretations and rationales (behind given courses of action) as constructed by those organisational members proximal to the decisions and events relevant for the study.

In conclusion, the above considerations have led to rejecting positivism and choosing

interpretivism as a more favourable philosophical perspective for investigating the organisational precursors to HAI issues.

Note that the selected paradigm, interpretivism, embeds a social constructionist assumption according to which the findings of organisational analysis of failure consist of interpretive constructions. Hence, with the choice of an interpretive paradigm it is easy to enter the realist vs constructionist debate, which notably put in opposition those assuming the objective existence of a reality 'out there' against those assuming that organisational reality is a product of social and cultural processes.

An article by Le Coze and Pettersen (Le Coze & Pettersen, 2008) has the merit of having presented this debate in a safety science attire. More precisely the author focused on the appropriate foundations for the area of resilience engineering (Hollnagel, Woods, & Leveson, 2006), which is a sub-area of safety science. In the article, on the one hand Pettersen suggests that critical realist ontology is appropriate given the layered conception of reality it embeds. *Critical realism* (Bhaskar, 1979) is one form of realism according to which observable regularities in the world can be explained by referring to the hidden mechanisms and structures that have caused them. This view integrates well with resilience engineering orientation "to shifts focus from actual events [observable at the system sharp end] to the underlying [organisational] facts and mechanisms...that are driving the functioning of socio-technical systems" (Le Coze & Pettersen, 2008). On the other hand, Le Coze reflects on the benefit that a social constructionist position can offer to resilience engineering. He warns that the validity of a model produced by a practitioner, should be assessed not in absolute terms but in relation to the viability of the model in the specific situation at hand. In this light, emerging theoretical models should be assessed against the experiences and background of those organisational members for which the model has been produced, and not by using some external seemingly objective criterion.

While both Le Coze and Petterson positions seem well grounded, they do not solve the dialectical tension between realism and social constructionism. It is possible to draw on Crotty to do so (Crotty, 1998). Crotty, has suggested that conflicts such as these stem from the tendency to equate realism to *objectivism*; however, the two are clearly distinct concepts: (i) the former is an ontological position assuming there is an external objective reality existing independently of the human mind; (ii) objectivism is instead an epistemological position asserting that meaning exists in objects independently of any human consciousness (Crotty, 1998). It follows from this that it is objectivism that should be opposed to constructionism, and not realism. In fact, stating that an external world exists independently of any human consciousness of it (realist assumption), does not rule out the possibility that meaning emerges out of the interaction between human

consciousness and the external world (constructionist assumption). In short, constructionism in epistemology can sit comfortably with realism in ontology (Crotty, 1998). It can be noted that the coexistence of realism (and his more articulated derivatives, such as critical realism described above) and constructionism is an accepted position within the area of sociology of risk (Beck, 2000, p. 212; Flynn, 2006, p. 86). Therefore, while these considerations downplay the importance of the realist vs constructivist debate, they also further reinforce the confidence on interpretivism as a plausible research paradigm for the present research.

APPENDIX C: EXAMPLE OF AN NTSB SAFETY RECOMENDATION LETTER

This appendix reports NTSB safety recommendation letter A-84-82 through -84. The safety recommendations conveyed can be found on the last two pages of the letter, i.e. pages seven and eight.

*Ng 1688*

NATIONAL TRANSPORTATION SAFETY BOARD  
WASHINGTON, D.C.

ISSUED: August 13, 1984

Forwarded to:

Honorable Donald D. Engen  
Administrator  
Federal Aviation Administration  
Washington, D.C. 20591

SAFETY RECOMMENDATION(S)

A-84-82 through -84

Between 4 p.m. and 5 p.m. on March 8, 1984, the National Transportation Safety Board received several telephone calls from witnesses who had observed aircraft flying close to tall buildings located in the Rosslyn, Virginia, area. These aircraft were conducting approaches to land at Washington National Airport, Washington, D.C. The witnesses were located on the ground and in the building at 1000 Wilson Boulevard. As a result of the reports and because of previous similar incidents investigated by the Safety Board, the Safety Board conducted a comprehensive investigation of the incidents. Ground witnesses, flightcrews, and air traffic controllers were interviewed, flight data recorders (FDR) from involved aircraft were read out, and recorded radar data were plotted. An analysis of this information has uncovered several safety hazards which warrant corrective action by the FAA. These involve the interpretation of descent profile altitude restrictions on instrument approach procedure charts, effectiveness of the minimum safe altitude warning system (MSAW), <sup>1/</sup> and air traffic controller procedures for issuing safety advisories to aircraft.

The official weather observation at Washington National Airport at 4 p.m. on March 8, 1984, was reported as ceiling measured 1,000 feet variable and broken, 2,000 feet overcast, visibility 6 miles with light rain and light snow; the wind was from 100° at 12 knots. The weather remained generally as reported at 4 p.m. until about 5:30 p.m. when, as the result of a frontal passage, ceilings lowered to indefinite, 200 feet, the sky was obscured, and visibility was one-quarter mile with thunderstorms and heavy snowshowers.

<sup>1/</sup> MSAW is designed to monitor aircraft with altitude transmitting equipment (Mode-C) for terrain clearance and to generate aural and visual alarms to controllers when an aircraft is at or predicted to be at an unsafe altitude.

-2-

The arriving flights were executing the very high frequency omni-directional range station (VOR)/distance measuring equipment (DME) standard instrument approach to runway 18 at Washington National Airport. The VOR/DME runway 18 approach is a nonprecision approach with a series of step-down descents beginning at an altitude of 3,000 feet at the 10-mile DME fix down to 900 feet at the 3-mile DME fix. After passing the 3-mile DME fix, an aircraft may descend to the minimum descent altitude (MDA) of 720 feet. The MDA must be maintained until the required visual references for the intended runway are identifiable to the pilot and the aircraft is in a position from which a descent to landing can be made at a normal rate of descent. The missed approach point is located at the 0.5 DME fix.

The Safety Board's laboratory analyzed the recorded radar data from the FAA's Automated Radar Terminal System III (ARTS III) to determine whether any of the flights deviated from the VOR/DME runway 18 approach procedures. The flightpaths of 21 approaches to National Airport between 4 p.m. and 5 p.m. were plotted using the aircraft course and altitudes to examine their relation to the standard approach procedure and the tall building located at 1000 Wilson Boulevard. Nine of the resultant flight profiles showed that the aircraft were descended below the 900-foot-altitude specified for the approach before reaching the 3-mile DME fix. Six of these flights were 200 feet low, one flight was 300 feet low, one was 400 feet low, and one was 500 feet low. An examination of the latter flight's flight data recorder showed that it had descended to 365 feet just after passing the 3-mile DME fix and before passing the tall building at 1000 Wilson Boulevard. The height of this building is 396 feet. At least two of these flight profiles indicated that the aircraft were flown dangerously close, either directly over or near abeam, to the tall building with no more than 100 feet vertical clearance. The Safety Board's investigation also revealed that nine MSAW alerts were activated at Washington National Approach Control and Tower, and that a controller took action on only two, to warn pilots of their low altitude.

#### Human Performance Aspects of Approach Charts

The Safety Board's investigation has focused on why the pilots of these nine flights deviated from the altitudes depicted on their approach charts. The crew of one of these flights stated that their interpretation of the 3-mile DME fix altitude of 900 feet was that it was only a recommended, and not a minimum, altitude. They further stated their belief that unless the word mandatory was depicted over the altitude, the flight was allowed to descend to the MDA altitude of 720 feet after it had passed the final approach fix (5 DME). Through discussions with other airline pilots, Safety Board investigators determined that many of them also had differing interpretations of what the 3-mile DME fix, 900-foot altitude meant on this VOR/DME runway 18 approach procedure. Some pilots viewed the altitude as recommended, some as minimum, and others as mandatory. This confusion suggests the need to clarify approach plates so that pilots understand the exact procedure to be followed.

-3-

The Safety Board compared the symbology used to depict altitude requirements on both the Jeppesen and the National Oceanic and Atmospheric Administration (NOAA)/Defense Mapping Agency (DMA) profile views of instrument approach charts, and found the two to be significantly different. The NOAA/DMA charts use the numerical altitude alone, or in conjunction with underscoring and overscoring lines, to depict different altitude requirements; for example, an altitude of 900 means "recommended altitude"; an altitude with an overscore, 900, means "maximum altitude"; an altitude with an underscore, 900, means "minimum altitude"; and an altitude with both an underscore and an overscore, 900, means "mandatory altitude." On the other hand, the Jeppesen charts use the numerical notation in conjunction with words to depict the various meanings. The altitude alone, 900', means "minimum altitude"; if other than "minimum altitude" is intended, the numerical notation is used in conjunction with the appropriate word above the altitude — mandatory, maximum, recommended.

900'            900'            900'

Many air carrier pilots were trained in the U.S. military service, which uses the NOAA/DMA approach charts. Also, many airline pilots are active in U.S. military reserve and national guard units, flying military aircraft on weekends and summer duty assignments. These same pilots, when flying for the airlines, generally use Jeppesen approach charts. Thus, these pilots use NOAA/DMA charts while flying for the military and Jeppesen charts while flying commercially.

The Safety Board believes that the lack of standardization of altitude legends on NOAA/DMA and Jeppesen charts could, in part, be responsible for the different interpretations of altitude requirements on approach charts. The Safety Board believes, therefore, that the same symbology or legend for depicting altitude restrictions should be used on both NOAA/DMA and Jeppesen charts to avoid misinterpretation by flightcrews.

This lack of standardization of chart legends is part of a bigger issue of serious concern to the Safety Board — that insufficient attention is given to human performance criteria in the review of approach procedure depiction on approach charts. These considerations go much farther than certain display requirements (e.g. width/height of letters) to include such items as amount of information displayed, and ease of identification and usability of that information. Pilots have been criticized for misinterpreting approach charts, but little consideration has been given to the operating environment in which the charts are used and the degree to which the charts themselves may be conducive to mistakes.

The Safety Board has investigated several accidents involving approach chart issues. In addition, testimony at public hearings has revealed that the FAA does not formally review approach charts designed by either the government or the private sector for human performance considerations. Further documentation of this concern is to be found in the NASA Aviation Safety Reporting System (ASRS) and the 1981 President's Task Force on Aircraft-Crew Complement. NASA has identified numerous ASRS incident reports in which approach chart issues played a significant role in the occurrence of the incident; and the President's Task Force emphasized that "the design and content of these charts should be improved."

-4-

In August 1982, the Safety Board issued a safety recommendation on this topic. In its November 1982 response, the FAA stated that it had undertaken a review of the subject area and expected this review to be completed in May 1983. The Board has not received the results of this review. Therefore, the Safety Board reiterates its previous recommendation that the FAA:

Establish formal human performance criteria for the development and evaluation of instrument approach procedures and instrument approach charts. (Class II, Priority Action)(A-82-91)

#### MSAW System

The second safety hazard of concern to the Safety Board as a result of its investigation of the March 8 incidents involves the MSAW system and its use by controllers.

Of the nine MSAW alerts activated at Washington Approach Control and Tower on the afternoon of March 8, 1984, a tower controller took action on only two to warn pilots of their low altitude. The Safety Board investigation focused on the reasons action was not taken on the other seven alerts. In order to verify that all nine MSAW alerts were properly transmitted to controllers, Safety Board investigators took the recorded air traffic control radar data to the FAA Technical Center in Atlantic City and ran the data on a "retrack program." Investigators witnessed the same radar and audio information that was presented to the controllers on March 8. The recorded radar data revealed that, in fact, all nine MSAW alerts were properly processed by the system to provide an aural alarm and video presentation in the tower.

When an aircraft descends below the minimum safe altitude a buzzer sounds over speakers in the tower control cab and in the radar room for about 5 seconds. The volume can be adjusted by controls located in both the tower and the radar room. The lowest possible volume setting is audible throughout the facility, but the system can be selectively inhibited or turned off completely. At Washington National Airport, MSAW alerts are heard simultaneously in both the approach control radar room and the tower cab.

An alert message also is visually displayed on the radarscope for the duration of the alert condition. The symbol "LA" appears above the alphanumeric identification tag associated with the affected aircraft. A separate area on the scope displays the abbreviation "LOW ALT" and a list identifying the aircraft which are causing the system to issue the low altitude alert. The system can list a maximum of five aircraft simultaneously which may be causing an alarm.

Aural and visual warnings also are used to alert controllers when aircraft come hazardously close to each other. This conflict alert system (CA) aural alarm is the same tone as that of the MSAW alert and the sound emanates from the same speaker. All controllers in the radar room and tower are able to hear the alarm when it activates. When a controller hears an alarm, he must examine his radarscope in order to distinguish between a low altitude alert and a conflict alert (which is displayed as "CA" above the aircraft data tag).

-5-

Safety Board investigators interviewed the two supervisors and four controllers who were controlling aircraft on the VOR/DME runway 18 approach during the period 4 p.m. - 5 p.m. on March 8. One tower controller remembered hearing two altitude alerts during this timeframe, and he issued the appropriate warning to the pilots of the affected aircraft. He did not recall hearing the other seven low altitude alerts. The other tower controller recalled hearing only one alert, but the two radar controllers and the two supervisors did not recall how many, if any, aural alerts were sounded.

The supervisor who was on duty in the radar room at the time explained that the conflict alert activates frequently when aircraft are on converging courses, even though they normally will be controlled so as to maintain prescribed separation criteria. According to the supervisor, "a great majority of the time you associate the aural tone with a conflict alert" and "over a period of years the thing goes off (CA) and you really don't pay attention to it 98 percent of the time." Unfortunately, when a controller ignores the aural alarm, he may be ignoring a warning from the MSAW system, rather than a conflict alert. It is a basic precept of psychology and human engineering that the ability of a stimulus to elicit a response (in this case, the ability of a warning tone to get the controller's attention) is reduced when the stimulus is habitually presented without a reinforcement. <sup>2/</sup> Reinforcement for a controller would be the acquisition of useful information from an aural alarm. In other words, when a controller is continually subjected to "nuisance alarms," i.e., those that are perceived as useless or distracting, he/she will pay progressively less attention to alarms. <sup>3/</sup>

The Safety Board found during its investigation of aircraft separation incidents at the Hartsfield-Atlanta International Airport on October 7, 1980, that the practice of ignoring alarms was prevalent. In its report, the Safety Board questioned the effectiveness of the CA and MSAW systems:

The frequency of conflict and low altitude alerts should be considered. This situation and the others mentioned above [common tone and source for LA and CA, and alarms which have their origin in another controller's airspace] results in repetitive alerts which, in turn, condition the controller to dismiss the alarms or alerts (i.e. the "cry wolf" syndrome). The Safety Board believes that improvements are needed in both the audio and visual cues for the low altitude and conflict alert systems. <sup>4/</sup>

<sup>2/</sup> Psychology and Human Performance, Robert M. Gagne and Edwin A. Fleishman, Holt, Reinhart and Winston, Inc., 1959, page 151.

<sup>3/</sup> Flight-Deck Automation: Promises and Problems, Earl L. Weiner and Renwich E. Curry, NASA TM 81206, June 1980, page 12.

<sup>4/</sup> Special Investigation Report: Aircraft Separation Incidents at Hartsfield-Atlanta International Airport, Atlanta, Georgia, October 7, 1980, (NTSB-SIR-81-6).



-6-

As a result of the investigation, the Board recommended that the Federal Aviation Administration:

Redesign the low altitude/conflict alert at ARTS III facilities so that the audio signal associated with the low altitude alert is readily distinguishable from that associated with the conflict alert and heard only by controllers immediately concerned with the involved aircraft. (Safety Recommendation A-81-134.)

In a letter dated December 21, 1981, in response to the Safety Board's recommendation, the FAA did not agree that separate aural alarms were needed, and stated, "we believe that the audio alarms represent a general warning or attention getter. The blinking alphanumeric represent the specific warning. It identifies the aircraft involved and the nature of the problem. The controller does not take control action based on the audio alarm; consequently, no benefit can be determined for the second audio alarm. The alarm or alarms mean the same thing, scan the display."

In a September 1, 1982, letter in further response to the Safety Board's urging that it adopt the recommended action, the FAA replied that it had "not changed our position" regarding A-81-134; therefore, the recommendation was classified as "Closed--Unacceptable Action" by the Board. In view of the information obtained during interviews with Washington National controllers during the investigation, the Safety Board believes that the FAA should reconsider its position on this matter and should implement action such as that called for in Safety Recommendation A-81-134 as soon as possible.

In a further attempt to determine why the local controller failed to issue low altitude alerts to seven of the nine flights which had caused the MSAW system to activate, the Safety Board examined the procedures for issuing safety advisories to airplanes. The FAA's Air Traffic Control Handbook, 7110.65C, paragraph 33, provides guidance to air traffic controllers to: "issue a safety advisory to an aircraft if you are aware the aircraft is at an altitude which, in your judgment, places it in unsafe proximity to terrain, obstruction, or other aircraft." Note 2 in paragraph 33 states in part, "recognition of situations of unsafe proximity may result from MSAW...." However, the Safety Board was given to understand that the phrase, "in your judgment," gives the controller the option, once the airplane has been identified on the BRITE display, to look at the airplane from the tower cab and form a judgment concerning the airplane's safety. If, in the controller's judgment, the airplane is a safe distance from obstructions and terrain, the controller may elect not to issue a low altitude alert.

During the investigation, Safety Board investigators discussed the provisions of paragraph 33 with senior FAA air traffic control management at Washington National Airport. Senior ATC management confirmed that a controller could decide not to issue a low altitude alert if, in his judgment, the airplane was at a safe altitude. They suggested further that this option may explain why the local controller may not have issued low altitude alerts to the pilots of at least some of the airplanes which had activated the MSAW system because of the option in paragraph 33.

-7-

In interviews with other controllers at Washington National Airport, the Safety Board received varied interpretations of the procedures contained in Paragraph 33. Some controllers stated that, in marginal weather conditions (ceilings close to the MDA), they would always issue a low altitude alert, while other controllers indicated that if they can see the airplane and it is not close to the buildings in Rosslyn, they do not issue the alert. The local controller on duty stated that he "didn't recall hearing the other seven MSAW alerts." The supervisor on duty in the tower, when asked how he judged whether an airplane was in unsafe proximity to the tall buildings, stated, "I can't tell what relation he is to that building when using the VOR DME approach.... I'm assuming he's flying the radial the way he's suppose to be."

The Safety Board is concerned that the provisions of paragraph 33 can lead a controller to nullifying the intent and objective of the MSAW system which is to alert a pilot when his airplane is at an unsafe altitude. The MSAW system software for approach path monitoring is programmed to activate if the airplane is 100 feet below the MDA, or if it is predicted that the airplane will be 200 feet below the MDA within 15 seconds. The Safety Board believes these activation parameters are definitive indications of unsafe proximity to terrain, and that the controller should not be called upon to make a judgment with regard to an airplane's safety. The controller should immediately inform a flightcrew of the activation of a low altitude alert, and any decisions concerning the airplane's safety should be made in the cockpit. The Safety Board thus concludes that the FAA should amend its procedures in paragraph 33 to eliminate the option apparently available to controllers to not issue a low altitude alert to an aircraft which has activated the MSAW system based on a visual judgment that the airplane is at a safe altitude and to require that a controller issue an alert to the flightcrew of all such aircraft.

Therefore, the National Transportation Safety Board recommends that the Federal Aviation Administration:

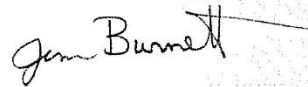
Prescribe standardized altitude symbology to be used in the profile view of approach procedure charts. (Class II, Priority Action)  
(A-84-82)

Redesign the low altitude/conflict alert at ARTS III/III-A facilities so that the audio signal associated with a low altitude alert is readily distinguishable from that associated with a conflict alert and so that it is heard only by controllers immediately concerned with the involved aircraft. (Class II, Priority Action) (A-84-83)

-8-

Amend the Air Traffic Control Handbook, 7110.65C, paragraph 33, to require a controller to issue immediately a low altitude alert to any airplane under his control which has activated the Minimum Safe Altitude Warning System. (Class II, Priority Action) (A-84-84)

BURNEIT, Chairman, GOLDMAN, Vice Chairman, and BURSLEY and GROSE, Members, concurred in these recommendations.

A handwritten signature in cursive script that reads "Jim Burnett". The signature is written in dark ink and is positioned above the typed name and title.

By: Jim Burnett  
Chairman

**APPENDIX D: STUDY 1 DATA SET**


---

**Table 29.** Complete data retrieved from the NTSB and FAA databases during the initial phase of the study (see § 4.4.1.1).

Safety Recommendation Letter	Safety Recommendation	Date Issued	Follow up exchange letters (identified by date)	
			FAA Letters	NTSB Letters
A-81-132 through -138 (2)	A-81-134	- Oct 6, 81	- Dec 21, 81 - Sept 1, 82	- July 8, 82 - March 8, 83
	A-81-135	- Oct 6, 81	- Dec 21, 81 - Sept 1, 82	- July 8, 82 - March 8, 83
A-83-27 through -30 (4)	A-83-30	- Mar 24, 83	- Jul 5, 83 - Dec 21, 83	- Nov 21, 83 - Feb 11, 85
A-84-82 through -84 (8)	A-84-83	- Aug 13, 84	- Jan 18, 85 - Aug 26, 85	- July 1, 85 - Jan 24, 86
	A-84-84	- Aug 13, 84	- Jan 18, 85 - Aug 26, 85	- July 1, 85 - Jan 24, 86
A-87-46 through -51 (4)	A-87-49	- May 22, 87	- Aug 25, 87 - Jul 12, 88	- Jan 22, 88 - Feb 21, 89
A-89-110 through -120 (12)	A-89-110	- Dec 5, 89	- Feb 23, 90 - May 4, 92	- Sept 26, 90 - Aug 31, 92
	A-89-111	- Dec 5, 89	- Feb 23, 90	- Sept 26, 90

## Appendix D

			- May 4, 92	- Aug 8, 92
A-90-124 through 132 (5)	A-90-130	- Sept 28, 90	- Dec 18, 90 - Oct 6, 93	- July 11, 91 - Jan 28, 94
A-90-16- through -163 (8)	A-90-161	- Oct 29, 90	- Jan 18, 91 - Dec 28, 92	- May 20, 91 - Apr 16, 93
	A-90-162	- Oct 29, 90	- Jan 18, 91 - Dec 28, 92 - Feb 3, 95 - Dec 15, 95	- May 20, 91 - Apr 16, 93 - Jun 20, 95 - Mar 21, 96
	A-90-163	- Oct 29, 90	- Jan 18, 91 - Feb 3, 95 - Dec 15, 95	- May 20, 91 - Jun 20, 95 - Mar 21, 96
A-94-186 through -188 (4)	A-94-186	- Nov 21, 94	- Jan 24, 95 - March 20, 95 - Sept 27, 95	- Apr 27, 95 - May 25, 95 - Nov 11, 95
	A-94-187	- Nov 21, 94	- Jan 24, 95 - Sept 27, 95 - Jan 26, 96	- Apr 27, 95 - May 25, 95 - Nov 20, 95 - Apr 8, 96
A-95-120 (3)	A-95-120	- Nov 30, 95	- Feb 21, 96 - Jun 19,96 - Jul 31, 97 - May 14, 98	- Apr 16, 96 - Jul 15, 96 - Dec 30, 97 - Oct 19, 98

Appendix D

			- Dec 13, 99	- Feb 3, 00
A-97-22 through -27  (5)	A-97-22	- Apr 16, 97	- Jul 1, 97	- Feb 27, 98
	A-97-23	- Apr 16, 97	- Jul 1, 97	- Feb 27, 98
	A-97-24	- Apr 16, 97	- Jul 1, 97 - Sept 25, 98	- Feb 27, 98 - Jan 14, 99
	A-97-25	- Apr 16, 97	- Jul 1, 97 - Sept 25, 98 - Aug 13, 99 - Mar 14, 00	- Feb 27, 98 - Jan 14, 99 - Nov 3, 99 - Jul 14, 00
	A-97-26	- Apr 16, 97	- Jul 1, 97 - Sept 25, 98	- Feb 27, 98 - Jan 14, 99
	A-97-27	- Apr 16, 97	- Jul 1, 97 - Sept 25, 98	- Feb 27, 98 - Jan 14, 99
A-06-44 through -47  (13)	A-06-44	- Jul 12, 06	- Oct 6, 06	- Sept 4, 07
	A-06-45	- Jul 12, 06	- Oct 6, 06 - Dec 29, 10	- Sept 4, 07 - March 14, 11
	A-06-46	- Jul 12, 06	- Oct 6, 06	- Sept 4, 07
	A-06-47	- Jul 12, 06	- Oct 6, 06	- Sept 4, 07
Totals				
11	25	-	54	56

## APPENDIX E: STUDY 2 INTERVIEW GUIDE

This appendix includes the interview guide used during Study 2 of the research (presented in chapter 6). The interviews are supposed to start from a general discussion on MSAW development and the questions are intended to support the interviewer in following that discussion—rather than being presented in a pure sequential and mechanical way.

Organisational Level	Questions
<b>Blunt end roles</b> (Senior/Middle Management/engineers/controllers involved in implementation)	<ul style="list-style-type: none"> <li>- Have you been involved in the MSAW implementation project? If yes, how?</li> <li>- Could you please give us a chronological overview of safety nets deployment in your organisation, with specific reference to the MSAW?</li> <li>- What was the rationale for adopting the MSAW?</li> <li>- What was/is the involvement of regulators in the introduction/management of the MSAW?</li> <li>- What have been/are the roles involved in MSAW development and management, how do they work, and how coordination between groups, and between groups and line management is established?</li> <li>- What lessons learned do you have for another ANPS about to enter the MSAW implementation process?</li> <li>- What is the policy to keep controllers in the loop?</li> </ul> <p>...any final comment you wish to make?</p> <p><i>Thank you for answering these questions.</i></p>

Operational controllers	<p><b>A. Tool Overview</b></p> <p>Could you please provide me with an overview of the tool?</p> <p>How does the tool help you in maintaining safety?</p> <p><b>B. Alert</b></p> <p>How do you realise there has been a tool alert?</p> <p>On average, how often does the alert goes off (e.g., once/twice per day/week...)?</p> <p>Considering the situations when the tool goes off, how many of these situations (indicate estimated percentage):</p> <p>you are already in control of the situation: _____</p> <p>you predict the alert to go off: _____</p> <p>the tool is indicating an imminent risk: _____</p> <p>Are there situations in which you would expect the tool to go off but it does not? (if yes please explain when and why)</p> <p>Could you inhibit the tool? (If yes please explain when and why?)</p> <p><b>C. Response Decision to Alert</b></p> <p>Could you describe me how do you decide how to respond to an alert?</p> <p>Are there situations in which you do not deliver the safety alert to the crew?</p> <p>How do you distinguish a relevant alert – i.e., an alert that correspond to a situation of real risk – from one which is not?</p> <p><b>C. Usefulness</b></p> <p>When is the tool most useful?</p>
-------------------------	--



	<p>When is the tool least useful?</p> <p><b>C. Training</b></p> <p>What training did you receive for operating this tool (duration)?</p> <p>What aspects do you feel were not covered by the training?</p> <p><b>F. Confidence (Trust) on the tool</b></p> <p>What are the positive characteristics of the tool that increase your confidence on it?</p> <p>What are the negative characteristics of the tool that reduce your confidence on it?</p> <p>What does not work so well, areas of improvement?</p> <p>...any final comment you wish to make?</p> <p><i>Thank you for answering these questions.</i></p>
--	--

## APPENDIX F: STUDY 3 QUESTIONNAIRE

---

### QUESTIONNAIRE

#### The Organisational Side of Safety Net Implementations

#### INTRODUCTION

---

- With this short questionnaire, we would like to collect your expert feedback over some organisational and managerial conditions that appear to influence the quality of the implementation and operation of safety nets. These conditions have been identified during a study sponsored by EUROCONTROL, and your expert opinion on them is needed to verify their validity and completeness.
- Each section of the questionnaire will present you a short definition of the identified conditions followed by two open questions that you are kindly invited to respond. There we invite you to describe examples from your professional experience whenever you feel it appropriate—we deeply value this kind of feedback.
- Please note that no name of people, company, airport or nation will appear on our reports. Confidentiality will be preserved.
- Upon completion, the questionnaire must be returned to: [simone.rozzi@gmail.com](mailto:simone.rozzi@gmail.com)

**PERSONAL INFO**

Name and Surname:	
E-mail:	
Organisation:	
Safety Nets Implemented in Your Organisations:	

1. What are you job title and your primary work task in your organisation?

--

2. Could you broadly describe your involvement in the domain of safety net?

--

**CONDITION A:           MANAGERIAL ASSUMPTIONS DRIVING ADOPTION AND  
IMPROVEMENT**

*Definition:* At the time of adoption, senior management might have over optimistic assumptions about the challenges related to the implementation of safety nets, about the ability of the manufacturer to implement these alarms, about the impact of safety nets on air traffic controllers' operational tasks, about the required expertise and air traffic controllers' involvement, about the implementation schedule.

A.1     Please comment the description above based on your direct or indirect experience (comments, reflections, examples from both your direct and indirect professional experience would be greatly appreciated).

A.2     Considering the condition we have just discussed, what recommendations would you give to senior managers willing to introduce safety nets in their ANSPs.

**CONDITION B: ORGANISATIONAL HANGLING CAPABILITIES**

Definition: The successful introduction of safety nets into the operation goes beyond their physical installation into the existing software infrastructure. It is also a learning process, where the ANSP develops focused capabilities to address the level of nuisance alerts and other challenges accompanying their introduction. Ultimately, the introduction of safety nets needs to be matched by appropriate support from senior management, availability of leadership, expertise, manpower, and supporting tools dedicated to the continuous improvement of safety nets.

B.1 Please comment the description above based on your direct or indirect experience (comments, reflections, examples from both your direct and indirect professional experience would be greatly appreciated).

B.2 What recommendations would you suggest to ANPS new to safety nets that wish to develop appropriate organisational capabilities in the domain of safety nets?

**CONDICTION C: CONTROL OVER IMPLEMENTATION QUALITY AT THE BOUNDARY BETWEEN THE SERVICE PROVIDER AND THE SOFTWARE MANUFACTURER**

Definition: When purchasing safety nets as built-in components of a larger ATM system purchase, the implementation of safety nets may be left entirely to the manufacturer. When this occurs, if requirements are not specified before purchasing the system, control over quality of the implementation and the tuning process can shift entirely to the manufacturer. This brings the risk that ultimately the final implementation reflects more the manufacturer viewpoint than that of the ANSP, potentially resulting in a less than optimal implementation. Equally, by adopting an approach that allows little or no specific expertise and skills about safety nets to be absorbed from the manufacturer, the ANSP has little opportunity to subsequently optimise the safety nets itself.

- C.1 Please comment the description above based on your direct or indirect experience (comments, reflections, examples from both your direct and indirect professional experience would be greatly appreciated).

- C2. What recommendations would you provide to ANPS new to the safety net domain in order to avoid the problem described above?

## CONCLUSION

- D1. Please describe here any relevant organisational and managerial condition important for the success of safety net implementation that has not been mentioned here.

- D2. Please feel free to express here any comment you may have about this questionnaire.

>>We sincerely thank you for having completed this questionnaire.

Please return it to: [simone.rozzi@gmail.com](mailto:simone.rozzi@gmail.com)<<

## APPENDIX G: LIST OF PUBLICATIONS PRODUCED FROM THIS RESEARCH

---

### Peer Reviewed Publications

- Amaldi, P., & **Rozzi**, S. (2012). Inter-Organizational Safety Debate: The Case of an Alarm System from the Air Traffic Control Domain. *International Journal of Sociotechnology and Knowledge Development*, 4(1), 30–47.
- Rozzi**, S. (2012a). A Longitudinal Systemic Framework for Identifying the Organizational Precursors to Flawed Human Automation Interaction in Safety Critical Domains. *Human-Machine Interaction (Formal H)*, 17.
- Rozzi**, S. (2012b). Applying the Resilience Engineering and Management Perspective to Problems of Human Alarm Interaction in ATM. Presented at the 2nd SESAR Innovation Days, DLR & Technical University Braunschweig, Braunschweig, Germany.
- Rozzi**, S., Amaldi, P., & Kirwan, B. (2010). IT innovation and its organizational conditions in safety critical domains: the case of the minimum safe altitude warning system (pp. 1B3–1B3).
- Rozzi**, S., Amaldi, P., & Kirwan, B. (2009a). Automated Safety Nets in Air Traffic Control: How underspecified policy of use and unexpected adaptation can hamper their intended benefit. Presented at the Close Calls Organizations, near misses, alarms, and early warnings, London School of Economics, London, UK.
- Rozzi**, S., Amaldi, P., & Kirwan, B. (2009b). Identifying how automation can lose its intended benefit along the development process: A research plan. Presented at the 9th Bi-annual International Conference on Naturalistic Decision Making (NDM9), BSC London, 23-26 June 2009.
- Rozzi**, S., Amaldi, P., & Fields, B., B. (2008). Task Analysis and Contextual Models of Controllers Activity for Interactive System Design: A literature review. Presented at the Innovative Research Workshop and Exhibition, EUROCONTROL Experimental Centre, Bretigny su Orge, France.



**Non-Peer Reviewed Publications**

**Rozzi, S.** (2013). Safety Critical Automation and its Organizational Precursors. Presented at the Summer Conference 2013, Middlesex University, London, UK.

**Rozzi, S.** (2013). Viewing Problematic Human Automation Interaction in Safety Critical Domains from an Organizational Safety Perspective. Human Work and Interaction Design, Ph.D course, Roskilde University, Denmark, May 15-17.

**Rozzi, S.** (2013). Developing a Safety Net Capability, from Ops Room to Senior Management. *EUROCONTROL NetAlert Newsletter*, (16 February Issue), 7–9.

**Rozzi, S., Amaldi, S., and Fields, B.** (2008) Information system evaluation, Unit of Analysis, and Perturbation Effects in Safety Critical Domains. Position Paper presented at the workshop CSCW and Human Factors: Where are we now and what are the challenges, Nov 8-12 2008, San Diego, CA.