



Analyzing Near-Miss Events: Risk Management in Incident Reporting and Investigation Systems

Carl Macrae

Analyzing Near-Miss Events: Risk Management in Incident Reporting and Investigation Systems

Carl Macrae

Contents

Abstract.....	1
Acknowledgements	1
Introduction	2
Airline incident reporting and investigation systems	2
Challenges of risk management practice	3
Research methods	4
Setting and participants	4
Data collection.....	4
Data analysis.....	6
Findings and discussion: Practices of managing organizational risk	6
Framework of risk assessment: organizational risk resilience.....	7
Implications: Resilience, defences and organizational control	11
Identifying risks: interpretive vigilance	12
Implications: Knowledge, ignorance and early warning signs	14
Resolving risks: participative networks	15
Implications: Organizing culture, learning and expertise	17
Broader applicability and limitations	19
Conclusions	19
References	21

The support of the Economic and Social Research Council (ESRC) is gratefully acknowledged. The work was part of the programme of the ESRC Centre for Analysis of Risk and Regulation.

Published by the Centre for Analysis of Risk and Regulation at the
London School of Economics and Political Science
Houghton Street
London WC2A 2AE
UK

© London School of Economics and Political Science, 2007

ISBN 978 0 85328 139 9

All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, without the prior permission in writing of the publisher, nor be otherwise circulated in any form of binding or cover other than that in which it is published and without a similar condition including this condition being imposed on the subsequent purchaser.

Printed and bound by Kube, December 2007

Analyzing Near-Miss Events: Risk Management in Incident Reporting and Investigation Systems

Carl Macrae¹

Abstract

This paper examines the practices that support an important risk management strategy in many organizations—analyzing and learning from past incidents of operational error and failure. Risk management has become an increasingly important managerial function, and a range of prescriptive standards and guidelines for risk management have been produced. The situated work practices that underlie risk management in complex organizations have, however, received little empirical attention. To address this gap, this paper presents a qualitative, inductive study of the operation of Incident Reporting and Investigation Systems (IRIS) in airlines. This research aimed to characterize and map the assumptions, beliefs, strategies and tactics that determine risk management practice in this setting. Practices of risk identification, risk assessment and risk resolution were examined and characterized through three innovative concepts: interpretive vigilance, organizational risk resilience and participative networks. Characterizing practice in this way allows a range of theoretical and practical implications to be developed concerning the place of organizational knowledge, control and culture in risk management.

Acknowledgements

This paper is based on my doctoral research undertaken at the University of East Anglia. This research was funded by the Economic and Social Research Council and a collaborative industry partner. I am grateful to the practitioners who gave up so much of their time during the course of this research. I would like to thank two anonymous referees and Bridget Hutter for helpful comments on early drafts of this paper, and my doctoral supervisors, Nick Pidgeon and Mike O'Leary, whose contribution to my thinking on these issues is greatly appreciated. Thanks also to Amy Greenwood for her helpful editing of the manuscript.

¹ Correspondence: Carl Macrae, Centre for Analysis of Risk and Regulation, London School of Economics and Political Science, Houghton Street, London WC2A 2AE. (c.macrae@lse.ac.uk).

Introduction

Incident reporting systems are used extensively in safety-critical industries such as commercial aviation and chemical processing to learn from near-miss events (Phimister et al., 2003). In these industries operational breakdowns can have catastrophic and far-reaching consequences. Personnel are encouraged to report minor failures or mishaps so that any underlying risks can be identified and responded to before they contribute to a serious accident (van der Schaaf, Lucas and Hale, 1991; Reason, 1997). Incident reporting systems are now being developed in a range of other risk management domains such as healthcare (Barach and Small, 2000) and banking (Muermann and Oktem, 2002). As such, understanding how near-miss incidents are actually used in risk management is both a pressing practical and theoretical problem.

This paper examines the practical work of investigators who manage incident reporting systems. It specifically focuses on risk management in commercial aviation, analyzing the work of airline safety investigators who manage flight safety incident reporting systems. The primary objective here is to characterize the beliefs, assumptions and tactics that support the analysis of near-miss incident reports in this setting. The paper draws on in-depth qualitative, ethnographic data on the routine work practices of airline investigators in order to conceptualize the core components of these risk management practices.

The paper presents three important findings. First, risk assessment predominantly involved assessments of the current organizational capacity to rebuff and respond to errors, rather than predictions of the possible future consequences of those errors. Second, risk identification involved actively searching out minor uncertainties and ambiguities in both current models of an organization's operations, and those operations themselves. Third, risks were resolved by creating temporary networks of operational specialists and experts, who could be widely distributed around the organization, to participate in the investigation and review of risks.

Airline incident reporting and investigation systems

Incident Reporting and Investigation Systems (IRIS) lie at the heart of many airlines' risk management efforts. These systems provide infrastructures that support analyzing and learning from past events: minor errors, failures and disruptions (Pidgeon and O'Leary, 2000). They also support the core tasks of risk management. IRIS are functionally rather simple. In airlines they are operated by safety investigators based in an internal safety oversight unit. These units have no executive authority, and are charged solely with monitoring and reporting on safety performance within the airline. Front-line personnel, such as engineers and pilots, are required to report any operational error or mishap they encounter that may have safety implications. These paper-based reports are entered into a dedicated management information system, and then reviewed and assessed by the investigators.

Investigators are responsible for identifying whether an incident represents a serious risk to flight safety, and for determining what the appropriate company response ought to be. This risk analysis process is predominantly based on professional judgement. Investigators typically have extensive experience to draw on, both of

operational work and of safety management. Most have lengthy first-hand experience as engineers or flight crew, and as safety engineers or management pilots.

Investigators can respond to any identified risk in a number of ways. Principally, investigators can request additional information regarding an event from the relevant specialist or manager. This may involve requesting an engineer check a piece of equipment, a manager debriefing members of crew, or a specialist reviewing a local procedure. Due to the limited information provided by the initial reports, any incident considered problematic has to be followed up in this way in order to make a fuller assessment.

From these initial enquiries, complex and lengthy investigations can result. These can involve a range of people both within the organization and outside it, such as ground service contractors, air traffic service providers and equipment manufacturers. These may last from a few days to several months, and in a large airline several hundred can be ongoing at any one time.

A second and equally important element of risk management practice in IRIS involves the communication of risks. Investigators report on safety events and distribute information on safety performance throughout the airline. Investigators prepare regular briefings to senior management that highlight recent significant incidents. They circulate regular newsletters to front-line personnel, reviewing noteworthy events, lessons learnt and action taken. And they have significant input into the safety agenda at board-level committee meetings, by preparing the focus and content of board papers for review.

Challenges of risk management practice

Current models of risk management vary widely, but all invariably focus on three core tasks: risk identification, risk assessment and risk resolution (e.g. IRM, 2002; Cabinet Office, 2002, COSO, 2004). These models specify that risks must first be identified by collecting and reviewing appropriate data. Then risks must be assessed to determine the level of threat they represent, typically in terms of the likelihood and severity of future consequences. Finally, risks must be resolved by evaluating each so that mitigating action can be prioritized on those that are least acceptable.

The functional simplicity of IRIS can belie the complexity of the challenges faced by investigators in practice. These challenges are particularly pronounced regarding these three core tasks of risk management: identifying, assessing and resolving risks.

First, identifying risks is hard. IRIS are awash with weak signs of potential risks. Investigators are presented with large numbers of incident reports—tens of thousands a year in large airlines. Any one of these may point to an important but previously unrecognized risk. Or they may not. Separating noise from signal is particularly challenging here.

Second, assessing risk is challenging. By definition, IRIS collect information on only minor and relatively routine disruptions that resulted in little or no actual harm. And airlines are complex and heavily defended systems. Multiple safety barriers and risk controls exist at every step. Estimates of future adverse consequences are therefore highly uncertain and often extremely low.

Third, responding to risks is problematic. IRIS are typically operated from internal but independent safety oversight units. Investigators have no executive authority and cannot mandate or enforce action. Their role is only to monitor, investigate and report on safety. Further, in airlines as in other complex organizations, risks can span several departments and specialist areas. Understanding and addressing these risks often requires the coordinated action of diverse teams of personnel.

This research therefore aimed to examine and explain the practical work that supports these three components of risk management in IRIS. It focused specifically on the practical assumptions, beliefs, models and tactics that underpinned risk management in airline flight safety incident reporting systems.

Research methods

A qualitative, grounded research approach was adopted to develop a theoretical account of risk management practice that both explained, and was well-grounded in, data on practice.

Setting and participants

The research was conducted in the safety oversight units of five airlines and two national air safety agencies. Data collection was primarily focused on the largest airline involved. A total of 26 investigators participated from these seven organizations, 10 in the main collaborating airline (Table 1). These organizations operated similar and well-established flight safety incident reporting programmes, though they differed in relative size and function. The airlines also differed in size. They included two large airlines, one medium-sized airline and two small airlines. The number of incident reports collected from pilots was broadly proportionate to their size (Table 1). The national air safety agencies collected reports from the entire national aviation industry, although these amounted to a similar quantity as the large airlines, as the agency's reporting requirements encompassed only a proportion of each airline's internally collected reports.

Data collection

Qualitative data on risk management practice was collected in three broad phases. This data collection was guided by the grounded theory principles of theoretical sampling (Glaser and Straus, 1967; Locke, 2001; Pidgeon and Henwood, 2003). The focus of data collection progressively shifted to examine practical issues and settings that provided the best opportunities to develop the emerging findings.

The first phase involved a range of interviews with investigators in the main collaborating airline (Table 2). It began with a set of unstructured familiarization interviews. These were followed by a set of semi-structured interviews that explored processes of expert decision-making by having investigators verbalize their assessments of ten typical incident reports.

Table 1. Summary of key participants and level of incident reporting across all participating organizations.

Organization	Participating investigators	Incidents handled per year (approx)
Principal collaborating airline	10	8,000
Other collaborating organizations		
Small franchise airline – A	1	200
Small franchise airline – B	1	500
Medium subsidiary airline	1	2,000
Large international airline	5	8,000
National regulatory agency	3	8,000
National safety agency	5	7,000
Total	16	–
Total	26	–

This phase was concluded by a set of semi-structured interviews that further explored the emerging themes by asking investigators to comment on a report summarizing the findings to date, and further focusing on the nature of investigators’ knowledge in risk management practice.

The second data collection phase involved three months of participant-observation (Lee, 1999; Robson, 2002) in the main collaborating airline. The daily risk management activities of investigators were studied through some 400 hours of observation. I sat with and questioned investigators as they assessed incident reports and observed a range of routine conversations and formal and informal meetings and tasks (Table 3). Ethnographic fieldnotes were taken by hand (Emerson, Fretz and Shaw, 1995) and later typed up, detailing the practices and discussions I observed.

Table 2. Summary of interview data sources.

Interview	Number of interviews
Unstructured familiarization	5
Critical incident protocol	7
Developmental thematic review	9
Semi-structured comparative	
Small franchise airline – A	1
Small franchise airline – B	1
Medium subsidiary airline	1
Large international airline	5
National regulatory agency	3
National safety agency	5
Conversational group review	2 (5, 3) ^a
Total	39

^a Numbers in brackets denote the number of participants in each group review interview.

The third phase of data collection began with semi-structured interviews at the additional six organizations (Table 2). These were selected on the basis of the similarity of their risk management systems, and their difference in terms of size and function. This allowed comparisons and contrasts to be drawn with the data already collected at the main collaborating airline. A final set of group review interviews were conducted at the main collaborating airline to get feedback on and develop the late-stage findings.

Table 3. Summary of key participant observation data sources.

Aspect of participant observation	Number
Weeks in the field	10
Incident reports talked-through	464
Weekly briefs discussed	20
Meetings observed	
Flight safety team meetings	5
Flight safety board meeting	1
Company operations review meeting	1
Franchise airline safety board meeting	1
Total	8

Data analysis

Data analysis followed the grounded theory principles of constant comparison (Glaser and Straus, 1967; Turner, 1983; Locke, 2001) and was ongoing throughout the research. The aim was to move from specific data instances and examples of practice to more general theoretical concepts and relationships. Interview and fieldnote data was transcribed. Then, all instances relevant to understanding risk management practice were labelled. These labels aimed to describe, in clear terms, that small aspect of risk management practice. Data analysis cycled from this initial and low-level coding to phases of higher-level integration, comparison and revision. Low level codes were synthesized into higher-level categories. These categories were gradually defined and related to one another, producing a conceptual account of the risk management practices observed. Iterative phases of low-level labelling and higher-level integration were conducted throughout the research until a coherent account of risk management practice was produced that fully captured and explained all relevant data instances. Where these data instances are presented throughout this paper, they are referenced according to the phase of data collection and the participant being quoted (e.g. Phase 2- Investigator 3).

Findings and discussion: Practices of managing organizational risk

The practices associated with risk identification, assessment and resolution in IRIS are characterized and explained here in terms of three core concepts: resilience, vigilance and participation (Table 4). These concepts aim to capture how risk management was conducted in practice. They characterize the assumptions, beliefs and tactics that risk management was based on in IRIS.

Table 4. Core concepts and their relation to practice.

Core concept	Aspect of practice
Organizational risk resilience	Framework for assessing risk and safety
Interpretive vigilance	Ways of identifying and interpreting risks
Participative networks	Means of addressing and acting on risks

This theoretical account of risk management practice provides an initial structure through which to examine and understand what is a relatively new empirical arena. As an exploratory, inductive analysis of practice, each of these concepts is

purposefully broad. They each draw on a broad heritage in the risk and safety management literature.

Ideas of resilience have long been associated with risk management (Wildavsky, 1988; Hood and Jones, 1996). Resilience typically refers to qualities of flexibility and adaptability in the face of surprising events: resilience allows individuals, organizations and societies to respond to, cope with and learn from adverse events (Sutcliffe and Vogus, 2003). Vigilance is a concept strongly linked to ideas of attention, alertness and error. It features in accounts of the attentive monitoring of failures that is found in highly-reliable organizations (Weick, Sutcliffe and Obstfeld, 1999). Its absence has been used in explanations of human error (Reason, 1990) and organizational disaster (Freudenburg, 1992). The participation of workers in safety management and risk regulation has become an increasingly salient feature of risk management programmes (Hutter, 2001; Wright and Lyons, 2001). Equally, new models of organizational learning emphasize the importance of participative, interactive processes in the production and use of knowledge (Lave and Wenger, 1991).

The concepts of resilience, vigilance and participation are therefore developed here as lenses onto a range of interconnected features and characteristics of organizational risk management. As such, they provide the basis for drawing on and developing a wide range of organizational theories in order to better understand risk management practice. Each concept is elaborated, detailing and explaining risk management practice in flight safety IRIS. The key theoretical implications of each are also examined.

Framework of risk assessment: organizational risk resilience

Assessments of risk were guided by a set of practical assumptions, principles and beliefs concerning the nature of organizational safety. Taken together, these constituted a working theory, or theory in practice, of risk. This framed the aspects of reported events that investigators attended to, and it determined how they interpreted and made sense of these incidents. This interpretive framework had two core components. First, risk and safety were interpreted in organizational terms, as properties emerging from organizational activities and processes—rather than in terms of consequences and outcomes as is the case in typical methods of risk assessment. Second, safety was viewed not simply as the absence of error and failure, nor as merely the ability to catch and contain mishaps. Safety was conceived of as the capacity to defend against the potential for minor mishaps developing further and escalating into more serious breakdowns. Attention focused on the adequacy of those organizational defences and controls that remained in reserve and unused in any situation—those that provided resilience to the risk of an event escalating. This view of safety can therefore be characterized as organizational risk resilience. Risk assessment in IRIS was used to determine where this capacity for resilience was eroded.

The core assumption that formed the foundation of risk assessment in IRIS was that the potential for catastrophic consequences always existed in operations. Human error and technical failure were considered inevitable and entirely normal features of organizational activity: all aspects of operations were “fallible and frangible” (Phase 1-Participant 5). IRIS collected reports of these operational failures every day.

Significantly, there was a pervasive and fundamental belief that major accidents could result from a unique combination of these otherwise routine and minor failures. One investigator explained this with reference to another airline's recent accident:

The classic one is... the causes of the Paris Concorde accident. And it's fascinating, the number of factors that led to that accident, fascinating. And any one of them would have stopped that accident, and yet they happen every single day—every one of them. (Phase 3-Investigator 5: P3-5).

The accident referred to here, the fatal Concorde accident of 25th July 2000, resulted from the aircraft running over a small strip of metal left on the runway during take-off. This burst a tyre on the left landing gear. Debris from the tyre damaged the underside of the wing and ruptured a fuel tank. Fire broke out as fuel streamed from the wing. The crew shut down engine number two on the fire warning. Engine one soon lost power—both affected by the fire engulfing the left wing—and the aircraft could not gain height. The right engines lost thrust seconds later and, unflyable, the aircraft crashed into a hotel.

The assumption derived from this accident, and every other, was that any minor event had the potential, however remote, to develop into a serious, unforeseen and catastrophic accident. As such, assessing risk on the basis of absolute 'worst case' potential outcomes, as many guidelines recommend, was considered impractical as "every incident would be a catastrophe" (P3-12). Instead, risk assessment in IRIS was directed at assessing the underlying organizational capacity to prevent minor mishaps escalating into major breakdowns. Simply put, when assessing incidents investigators were literally interested in:

What are the things that stop this being a catastrophe? (P1-2).

Investigators were concerned with how resilient operations were—how effective risk controls and safety defences were at correcting or containing small failures. What was of primary importance in these assessments was not simply how well an error or failure had been caught by controls, but what defensive capacity remained beyond those, in case the event had developed further.

A key principle underlying risk assessment here was that numerous diverse defences should remain in any situation, protecting operations against the potential for an event to escalate further. That is, operational safety required not simply resilience to failure, but resilience to risk. One example that highlights this approach was a hypothetical 'worst case' scenario in which a flight crew receive a warning from the Ground Proximity Warning System to urgently pull up away from terrain:

The GPWS is the classic really, that's a worst-case, if you get a real one you've lost it really, it's one circuit breaker left to protect you. You are down to tight stuff. (P1-1).

While a safety system might do as it was designed to and catch the problem, situations such as these were nonetheless considered "as bad as it gets" (P2-1). With no further defences remaining, a failure here would leave operations entirely unprotected and exposed to potential catastrophe.

In light of this, incidents were used to assess the underlying organizational capacity for safety—characterized here as risk resilience—rather than to assess the actual adverse impact of an incident or to predict its potential future consequences.

Assessments of the quality of organizational safety fell along a broad spectrum. They tended to fall into three key distinguishable groups. First, where defensive processes

were deemed adequate and currently acceptable—for the time being. This broadly represented a ‘normal’ and unproblematic state, and a close approximation to relative ‘safety’. Second, where risk resilience was considered to be reduced, and existed in a somewhat deteriorated state. This was where defensive processes were considered less robust or extensive than they could be, or were functioning below par. Operations remain protected, but in a less than optimal condition. Third, where risk resilience was entirely degraded, deficient or absent. Here, a state was deemed to exist where defences were deficient, providing few or only weak protections against minor failures escalating into major ones.

Assessments of risk drew together a range of subtle judgements that investigators made regarding the efficacy and adequacy of safety defences and risk controls. These judgements primarily focused on the organizational nature of defences. Safety defences or risk controls were viewed as properties of practical work routines that provided capabilities to deal with failure. Work routines involved the interaction of social and technical elements: people operating, managing and maintaining technologies. That is, defences and controls were not viewed as simple mechanisms, or purely technical objects, but as organizational processes that hinged on the cognizance and competence of personnel in relation to technical systems.

One investigator’s discussion of the onboard warning system ‘TCAS’ (Traffic Alert and Collision Avoidance System), which alerts pilots to aircraft on a course conflicting with their own, indicates the inherently social nature of even this apparently automated, technical ‘defence’:

You want to see the crew saying, we are likely to get a [traffic] warning out of this, and you want ATC [Air Traffic Control] to pick it up too. So that is a system getting close to the edge but the monitoring is good, so when the RA [traffic warning] does go off, the crew is ready for it. So that’s a system that even though a defence operated, it didn’t operate as the final, ‘holy shit what was that?’ So that is what counts, multiple systems and the crew at least having some knowledge of what is going on. (P1-1).

Investigators didn’t draw a neat line between technical and human aspects of the organization. When assessing risk, both were seen as complementary and inextricably intertwined. Defensive capabilities were analyzed as interactions between social and technical aspects of practical work—even though, for practical ease, a ‘defence’ was typically referred to as a single and static ‘thing’.

Investigators’ understandings of the Ground Proximity Warning System, or ‘GPWS’, which automatically warns flight crew of terrain hazards equally demonstrates this point. Investigators often referred to the GPWS as simply ‘a defence’. But what they were referring to, as revealed by their detailed discussions and in their practical examination of GPWS incidents, was actually an inherently sociotechnical process, rather than a single mechanism that is either in place and working or not.

Investigators believed that, on its own, the warning system itself only provides a warning. Flight crews have to act on that warning appropriately and competently for it to do any good. For that to happen, appropriate training is required and suitable procedures for responding to these warnings must be developed, implemented and monitored; So investigators understood each ‘defence’ as an admixture of social and technical processes that extended throughout the organization, and that worked in interaction. What was labeled a ‘safety defence’ or ‘risk control’ was, in practice, less static and determinate than may typically be assumed.

In assessing the effectiveness and quality of defensive processes, six attributes of organizational processes were routinely attended to (Table 5). These organizational attributes were core properties of risk resilience, and perceived deficiencies in these attributes were interpreted as signs of degrading risk resilience.

Table 5. Properties of organizational processes attended to in risk assessment.

Characteristics of organizational resilience and illustrative examples

Ability – the provision of appropriate competence, knowledge, skills, conditions, tools and equipment.

Example 1. A flight crew improperly used a new automated system.

A crew reported that while preparing for take-off they used a newly introduced automated system to calculate the airspeed, but then erroneously reverted to the old routine of setting flaps to '15' out of habit. Investigators were concerned that the introduction of this new system had created a performance trap ("this is a trap that they've obviously fallen into... they let experience override it," P4-6) that challenged crews' ability to perform this task effectively.

Awareness – the maintenance of requisite understanding, cognizance, comprehension, models, information and data.

Example 2. Engineers initially missed the cause of an electrical fault.

Passengers reported an electrical smell that dissipated once the In-Flight Entertainment (IFE) equipment was isolated. Engineers concluded debris around the wires was to blame, cleared it and reinstated the IFE for the return flight. More detailed investigation later revealed that screws had been driven through the wiring. "They convinced themselves into thinking it was just the debris, so it wasn't investigated properly down route... Then later they found the real problem" (P4-6). These fault finding processes had failed in terms of awareness.

Communication – the passing of timely and appropriate information supporting mutual comprehension and co-ordinated activity.

Example 3. Flight crew did not communicate effectively during a go-around.

An incident where a go-around was inadvertently initiated and then poorly flown, elaborated in Example 10 (Table 6), largely implicated poor communication between the flight crew. The crew were slow to share their understanding of the situation with each other, "the problem was they had two different mental models of what they were doing" (P4-6), and so literally "they weren't working together as a crew" (P4-3). Communication and coordination had broken down.

Verification – the means to check, substantiate, observe, monitor, record and confirm.

Example 4. Take-off was initiated at too low an airspeed.

Aircraft rotation (pulling the nose up off the runway) was attempted before the necessary air speed had been reached, due to a misplaced 'speed bug' marking the take-off reference speed on the air speed indicator. The missed bug hadn't been caught by the other pilot as they were awkward to check: "It's hard to cross-check from the right hand seat due to the angle of the dials" (P4-4). This reduced the efficacy of cross-check practices on this aircraft type.

Specification – the creation, use and suitability of procedures, plans, controls, customs, policy and conventions.

Example 5. Loose bolts caused a spurious engine fire warning.

Some bolts were not tightened at the end of a maintenance job. The initial stage of the job was performed by a new engineer as laid out in the manual, loosening all bolts, but the local convention was only to remove a few that were necessary for access. The engineer who completed the job then only refitted the bolts typical for local practice. Work practices were not properly specified. "The problem... is that if it says all those things on the card then they should do them all. So the customs should be reviewed and woven into the [manual]" (P4-5).

Margin – the maintenance of redundancies, reserves, alternates, excesses and buffers beyond the minimum required

Example 6. Flight crew landed with less than the required reserve fuel left.

A crew used some of their emergency reserve fuel while waiting for a snow storm to pass before landing. "They had forty minutes of hold fuel... the hold turned out to be forty two minutes, then a heavy snow storm came out of nowhere—well they just don't come out of the blue! ...that's poor management and planning, it's all going to plan but they're not looking far enough ahead" (P4-1). This indicated that appropriate reserves and margins were not being preserved.

Implications: Resilience, defences and organizational control

At core, practices of risk assessment in IRIS involved assessing the quality and extent of organizational control. Importantly, risk was not routinely assessed by attempting to predict the possible future consequences of specific incidents, as models of risk management typically specify (e.g. IRM, 2002). Instead, risk assessment was more of a diagnostic activity. It involved diagnosing where organizational processes did not provide adequate protection against failures, and then assessing the severity of this degradation.

These findings are in line with a range of research that increasingly focuses on understanding resilience in organizations that operate in unforgiving environments (e.g. Reason, 1997; Pidgeon, 1998; Weick, Sutcliffe and Obstfeld, 1999;). The concept of resilience has attracted attention in terms of the ways organizations can respond to and recover from unexpected disruptions (Hollnagel, Woods and Leveson, 2006; Sheffi, 2005).

At core, resilience has previously been conceived of as the ability to “bounce back” (Wildavsky, 1988, p. 77; Weick and Sutcliffe, 2001, p. 14) and learn from errors and failures as they occur. That is, resilience has traditionally been a concept premised on reacting in hindsight to failures and challenges once they are underway (Sutcliffe and Vogus, 2003). This strategy of resilience is typically held in opposition to that of anticipation, or foresight, which aims to predict and then put in place protective measures to guard against future threats—the premise that underpins modern risk management systems.

In practice here, however, ideas of resilience were heavily integrated with notions of risk. Here organizational safety was construed not only as this ability to ‘bounce back’ from actual events as they occurred. Assessments of safety went beyond this. Investigators assessed operational incidents in terms of the quality of the immediate organizational response to the event, but also crucially in terms of the residual, remaining systems that could have caught the event if it had progressed further, and the initial responses had failed. Safety was premised on the widespread assumption of the need for ‘defences in depth’ (Reason, 1990). That is, in any situation, investigators assessed the quality of the remaining defences or latent controls that were not called upon in the event, but nonetheless provided additional means to catch, control or correct errors if they had developed that far—such as multiple back-up systems, additional cross-checks of data or automated alarm systems. In normal operations, these systems are rarely used to ‘catch’ or ‘bounce back’ from an error. But they are nonetheless routinely maintained, monitored or performed to provide this capacity for resilience in case they might be needed. It was these ‘reserve defences’ that were the primary focus of investigators’ assessments. And safety was considered acceptable only when the potential for an error developing further or getting worse was adequately defended against by these reserve systems—in short, where there was resilience to risk in operations.

This study of risk management practice provides a partial sketch of how concepts of risk and resilience may be integrated. It also implies that research on organizational resilience should more closely examine the question ‘resilience to what?’—actual errors and events, possible failures and mishaps, or the risk of minor incidents escalating into catastrophe? Current research largely leaves this question open.

One of the most notable developments of existing theory from this study of practice centres on the ideas of safety defences, barriers and risk controls (Hollnagel, 2004; Reason, 1997; Svenson, 1991). Studying practice suggests that the way safety defences, risk controls—and operational hazards—are currently defined may be somewhat problematic.

In most accounts it is assumed that defences are some form of barrier designed to keep ‘bad’ things ‘out’. Like the concentric fortifications of a castle, safety barriers stop hazards—fire say, or explosions—coming into contact with valuable and fragile assets, such as people or property (Reason, 1997). These are the assumptions that ‘barrier’ theories of safety (Hollnagel, 2004; Svenson, 1991), as well as basic models of risk management, are based on. Accidents happen when hazards penetrate system defences and cause damage. But in practice, defences were not viewed as a barrier that separates an asset from a hazard. In fact, the traditional notion of ‘hazard’ barely featured. Threats were viewed instead as the possibility of organizational processes breaking down, and as weaknesses in the capability for effective control. That is, defences were conceived of as keeping organizational activity within safe, known and manageable confines. Defences didn’t stop a hazard getting ‘in’. They stopped operations getting ‘out’—out of control and out of the normal, safe range of activity (e.g. Reason and Hobbs, 2003). The risk was organizational processes themselves becoming unwieldy or fragile. This represents a subtle but important shift in emphasis regarding the nature of safety defences and risk controls, in that they may be better defined in theory and assessed in practice in terms of the degree to which they constitute and control effective organizational processes.

Identifying risks: interpretive vigilance

Risk identification concerns how risks come to be noticed and initially known about in organizations. Practices of risk identification in IRIS aimed to create a high degree of alertness and sensitivity to early, weak signs of previously unknown risks (Macrae, 2007). In light of their conception of risk, investigators specifically worked to actively seek out deficiencies in organizational safety: where there was little left to prevent errors potentially spiralling out of control. This approach to risk identification, and the associated analytical practices, can be characterized as interpretive vigilance. This encompassed two interrelated processes. First, investigators aimed to monitor and interpret incidents vigilantly for signs of potential operational problems. Second, they sought to remain vigilant to flaws in those interpretations and inadequacies in current knowledge of operations. As such, investigators strived for similar ideals of oversight as have been observed in other high-reliability organizations: where attention is focused on small operational deviations and disruptions (Rochlin, 1989), and where unexpected or surprising events become the locus for further enquiry and change (Weick and Sutcliffe, 2001).

A distinct set of assumptions and beliefs underlay risk identification in IRIS. Investigators viewed their role of risk management and oversight as one of maintaining knowledge and awareness of all risks facing the organization. A core assumption—and aspiration—here was that signs of new emerging risks could and should be identified by piecing together cues in apparently inconsequential, minor, ‘small’ events. Discovering that a problem had not been noticed or fully understood for some period of time was considered a serious failure, both of organizational safety and their own ability to monitor it:

The ones that trouble you are the ones that come out of leftfield and you think, how did we get caught out like that? ... You've not done your job by flagging this back earlier. (P1-1).

A core belief amongst IRIS investigators was that their job was to identify risks 'early' to avoid being 'caught out' and surprised by problems that had gone unnoticed for some time. Yet, while being caught out was considered an unacceptable failure on their part ("it means we have been derelict", P1-8), it was also considered inevitable. Investigators frequently discussed their own experiences of missing and misinterpreting early signs of risks. Knowledge of risk was assumed to be always partial and incomplete. Dealing with a continuous stream of incident reports reinforced this assumption:

We do our hazard identification and risk analysis and think we've cracked it.

And then things happen: events. (P3-6).

IRIS continually exposed investigators to information about unexpected and previously unrecognized organizational problems. Because of this, interpretations of risk were handled in IRIS with a degree of humility and caution. Risk assessments were seen as a product of current knowledge and understanding, and so continually open to change: "your judgement of risk is dependent on what's happened in the past, it's not a fixed entity" (P2-2). In IRIS it was widely acknowledged that at different times near-identical incidents could be interpreted in very different ways, "because your knowledge has moved—and hopefully expanded" (P2-3).

This set of fundamental beliefs underpinned the risk identification strategy used in IRIS. This strategy largely involved monitoring for potential gaps or inconsistencies in current knowledge of operations. Investigators aimed to identify where their knowledge of organizational safety was in some way questionable or suspect, to focus further investigation there. While this strategy was relatively straightforward, putting it into practice was not. In IRIS, identifying new or previously unrecognized risks was a task that, at the outset, was based on extremely limited data—often two or three events at most. Tracking trends of events was a useful indication of underlying problems. But, as investigators frequently noted when trying to identify risks in their early stages of development:

...you can't back it up with enough data. The bigger the trend the better, but the [IRIS] database won't always tell you that, because it is three examples that lead you to a conclusion. (P3-5).

Risk identification was therefore largely an interpretive rather than a statistical exercise. Incident reports were used to actively interrogate and test current models and understandings of operations, to find where these appeared weak or out of date.

Four distinct interpretive tactics were used in IRIS to make sense of incidents and determine where operational safety—and knowledge of it—may be weaker than previously believed (Table 6). Connections were drawn between minor incidents and past major accidents or current concerns, indicating that operations were exposed to a previously unrecognized threat. Patterns were made between similar features in events, indicating that a common underlying problem might exist. Discrepancies were sought out in operational processes or in current models of them, indicating that these needed re-examining. And novel facets of failure were focused on, indicating that current knowledge was incomplete in some way.

Through these interpretive tactics, the adequacy of operational safety and organizational knowledge were continually tested. IRIS provided a platform for identifying risks through the vigilant monitoring and interpretation of minor organizational mishaps.

Table 6. Illustrations of the interpretive tactics used in risk identification.

Interpretive tactics and illustrative examples

Drawing connections

Example 7. An aircraft nearly used the full length of a runway to land.

The crew reported that on an approach in heavy rain they failed to override the automatic reverse thrust due to unrelated confusion over the apparent failure of a windscreen wiper. This event was immediately deemed “a bit of a QF1” (P4-1), referring to the flight code of another airline’s aircraft that had overrun a runway and ended up in a field a few years previously. In that case, a water logged runway, poor crew communication and an inadequate braking technique were contributory factors. These factors were the basis of the connection drawn between that accident and this incident. This connection led the investigator to suspect that a superficially inconsequential incident may point to an emerging and unrecognized problem in landing and approach discipline.

Making patterns

Example 8. A series of events indicating improperly secured cargo.

Over a couple months, investigators noticed several similar events involving pieces of cargo in the aircraft hold being found, on arrival, to be improperly fastened down, not fastened at all, or flight crew reporting hearing a ‘thump’ or ‘bump’ during flight. Unrestrained cargo can be a problem if it moves around and affects the trim and handling of the aircraft. Three events had been reported in the first month, and then it went up to about seven in the second month, and they had been seeing “bigger lumps of cargo” moving around (P4-4) into the bargain. Although the incidents themselves had little actual impact, investigators flagged this up as a “minor issue snowballing” (P4-4). These events presented a clear pattern, suggesting that something was amiss in the loading of cargo. On closer examination, investigators found that all the events could be traced back to the same terminal, reinforcing and localizing their suspicion of an underlying problem with work practices there.

Sensing discrepancy

Example 9. An engine warning lead to a cancelled take-off.

A report described how a take-off had been aborted at low speed due to an engine overspeed warning that signalled one of the compressor fans in the jet engine was spinning too fast, reducing thrust. The flight crew contacted the engineering control centre, were advised to check the engine with two stationary engine runs and, as they were clear, to depart as planned. Rejected take-offs are part of normal operations. But on this occasion the investigator “thought they would have done other checks before restarting” (P4-3) with that particular type of warning and was not sure that the advice was appropriate on that occasion. Following this up the powerplant engineers confirmed that it was a ‘red’ EICAS (Engine Indicating and Crew Alerting System) warning that required additional precautionary maintenance checks before dispatch. This led to the advisory performance and procedures of engineering control being called into question and subjected to further investigation.

Perceiving novelty

Example 10. A switch misselection on approach was poorly responded to.

A crew reported they had inadvertently hit the ‘Take-Off/Go-Around’ (TO/GA) switch—which applies full engine power—instead of the autothrottle disconnect switch during approaches to land. This error was well-known as the switches are located in opposite positions on one type of aircraft compared to all the others built by the same manufacturer. This had been extensively investigated and was viewed as “design-induced pilot error” (P4-1) that could only be addressed through training when crew members moved to that aircraft type. However, in this event the crew were slow to realize what had happened, did not cancel the switch, and subsequently completed an awkward and clumsy go-around manoeuvre. This highlighted “a whole bunch of issues beyond TO/GA and the physical selection of the switch” (P4-6) such as crew communication, go-around training regimes in the flight simulators, and the effectiveness of crew briefings. This novel event overturned previously accepted beliefs in this area: “we thought selecting TO/GA was not in itself an unsafe action. This altered our view” (P4-6)

Implications: Knowledge, ignorance and early warning signs

Practices of risk identification in IRIS represent an approach to risk management that was based on a deep appreciation of the limits to knowledge, and that employed interpretive tactics that aimed to find gaps and inadequacies in current understandings of risk. Primarily, this indicates a core but relatively under-explored theoretical issue—that risk management is crucially concerned with acquiring and maintaining knowledge of the challenges and threats that face an organization. Managing risk is a process of developing and revising current knowledge in the face of irreducible uncertainty (Wildavsky, 1988). Models of risk management typically specify the infrastructures of information collection and manipulation that can

support this. But studying risk management practice highlights the intuitive, generative, creative—and often fallible—sociocognitive processes that seem essential to interpreting and making sense of ambiguous risk events in organizations.

Identifying risks was largely an interpretive rather than statistical process. It involved piecing together the limited, ambiguous and fragile information available on incidents with broader frames of reference and bodies of knowledge (e.g. Weick, 1995a; Cook and Brown, 1999). Viewing risk identification as an interpretive process in this way leads to several important insights.

First, the available knowledge and frames of reference are as important as the information and data collected. While most effort is focused on the collection of relevant risk data, this may be only half of the story. Here, detailed knowledge of operational processes, organizational goals, previous problems and past accidents were of particular importance in identifying new risks. How such knowledge is used, shared and represented in risk management—often in the form of stories about past events (e.g. Orr, 1996; Weick, 1987)—is a critical issue that requires closer attention. Further, these findings reinforce the argument that these forms of detailed, practical knowledge may easily be lost in efforts to systematize and standardize risk management processes (Power, 2004).

A second important implication is that ‘warning signs’ of organizational problems have to be actively constructed and created. They are not self-evident, to be merely noticed or missed by managers (cf. Kiesler and Sproull, 1982; Sheaffer, Richardson and Rosenblatt, 1998). The processes involved in interpreting warning signs is central to understanding risk management practice in organizations (Vaughan, 1996; Turner, 1978). In practice here, warning signs were constructed by forming a tentative belief that current knowledge was in some way inadequate or incomplete. That is, risk management was an interpretive process of testing and identifying the limits of current knowledge. Incidents were used to interrogate the unknown (Wildavsky, 1988; Macrae, 2006), test assumptions (Turner and Pidgeon, 1997), challenge expectations (Weick and Sutcliffe, 2001) and so become aware of ignorance (Smithson, 1989).

Practices of risk identification were organized around finding areas of organizational ignorance—where understandings of operational processes were poor, incomplete or confused. This broader consideration of ignorance and the limits of knowledge seems central to explaining risk management practice, in contrast to the narrow definition of uncertainty that is typically deployed in risk management models that focus exclusively on the probability of future adverse consequences. Further exploring the nature and purpose of different forms of ignorance (e.g. Smithson, 1989; Cunha, Clegg and Kamoche, 2006) in risk management practice would seem a valuable research endeavour.

Resolving risks: participative networks

IRIS supported a distributed and decentralized approach to acting on and resolving risks. As an oversight mechanism, IRIS did not provide direct authority to mandate or enforce action. Instead, the aspirations and practices of investigators in IRIS centred on managing awareness of risks within the organization, and creating ownership and accountability for resolving risks by relevant personnel. IRIS were

used to create participative networks around risks. Investigators initiated, coordinated and monitored the participation of personnel from around the organization—and often outside it—in examining and resolving risks. Participation is a key goal in many systems of risk management (Hutter, 2001; Wright and Lyons, 2001), and in models of organizational learning (Lave and Wenger, 1991). Here, risks were addressed by guiding local specialists to review and investigate operational practices in their area. In this way, investigators supervised the local and decentralized investigation of risks, and acted as the coordinating hub of often widely dispersed investigative activity. By organizing local participation in risk management, investigators were able to shape the production of organizational knowledge, change local work practices, and orchestrate organization-wide learning in response to risks.

A central belief of investigators was that the ultimate purpose of IRIS was to generate action and resolve risks. Incident reports were used to provoke and trigger deeper investigation of risks, to focus organizational resources on problems and, essentially, to effect organizational change. Assessments of risk were used to drive action and get things done, as investigators commonly explained:

The purpose you do [risk assessment] is... because you want to understand the appropriate corporate response. It helps you focus your resources on the appropriate areas to bring about the change that will hopefully get rid of that problem. (P1-7).

While organizational action was the aim, IRIS provided no direct control or authority over this action. In IRIS, creating organizational action was about influencing and guiding the action of others. In practice, investigators worked to achieve this in two ways. First, they publicized signs of potential risks within the organization. Directors, managers, supervisors and operational specialists were informed of significant incidents by the distribution of regular—typically weekly—briefings. These were intended to make sure that people who should know about serious incidents, did.

It means it gets to people. With the best will in the world, the director wouldn't have time to go in to [the IRIS database] to search himself, and so someone would be doing it for him. So if they aren't telling him the good stuff, then we are. (P2-1).

Second, investigators posed questions about the adequacy of organizational processes, by requesting reviews and further information from operational specialists, and by producing regular papers for review at board level.

In posing questions and publicizing incidents, investigators aimed to direct attention and manage awareness of risks within the organization. The aim was straightforward: to “flag it up, and say these are the things you should be looking at today” (P3-1). Through the IRIS, investigators were setting and continually revising a risk agenda—what people should be aware of and attending to. Equally, investigators aimed to provoke enquiry and action—to make risks visible to those responsible for resolving them:

[We] don't take the decisions, we invite the departments to—to take ownership of safety... It is making them accountable, and it is their decision. (P3-6).

By creating awareness of current risks within the organization—and an audit trail—IRIS were used to make the ownership of risks visible, as much as the risks themselves. Investigators aimed to ensure widespread participation in risk management. By distributing signs and questions regarding risks, IRIS were used to engage networks of diverse specialists and personnel in the investigation and

management of risks. Investigators contacted personnel who were relevant to resolving a risk, often bringing together a range of people from different operational areas, forming a temporary and distributed team. Incident reports provided a focus around which reflection and action was organized.

The extent of intervention required by investigators in resolving a risk was determined according to three criteria. These criteria were used to evaluate how worthwhile and beneficial it would be to act on a risk—in simple terms, the risk management value of action. Judgements of the value of acting on a risk considered three issues. First, the potential safety benefits that would accrue, in terms of strengthened resilience and greater knowledge. Second, the extent to which current organizational activity was already addressing the risk. And third, the practicality of actually addressing the risk and achieving material improvements (Table 7). These criteria were the basis on which the level of appropriate action was evaluated in IRIS.

Table 7. Practical criteria for evaluating the worth of risk management intervention.

Criteria for evaluating action and illustrative examples
<p>Potential learning – the extent to which safety may be improved by acting, in terms of strengthened risk resilience and increased organizational knowledge.</p> <p><i>Example 11. Unknown chemical powder spilt in the cargo hold.</i></p> <p>When the chemical was found, “they tested it for anthrax and everything, but did it damage the aeroplane? ... we’re concerned about the effect on flight safety—and that no one else seems to have worried about that” (P4-4). This was identified as a simple and direct opportunity to act to both specifically highlight the safety issues surrounding potential corrosive substances and aluminium airframes, and more generally to increase awareness of safety issues and increase the focus on safety.</p>
<p>Current activity – the extent to which current organizational activity is already addressing a risk, and how adequate and effective this action is</p> <p><i>Example 12. A series of events with aircraft marshalled too close to a building.</i></p> <p>Marshallers as one airport repeatedly directed aircraft extremely close to a building, resulting on one occasion with an engine striking a concrete pillar. While initially it was believed that an adequate plan had been drawn up by their counterparts at the foreign station (“these guys are on top of it” (P4-4), when more events occurred several weeks later, it became clear that these actions were ineffective and it was worthwhile them intervening: “[The station managers] are saying, well we’ve done everything and briefed them—well they obviously haven’t as it still keeps happening!” (P4-5).</p>
<p>Pragmatic practicality – the extent to which it is practically possible for the organization to address a risk and achieve material improvements in safety</p> <p><i>Example 13. Flight crew seat moved on fastening rails during landing.</i></p> <p>A pilot reported that while landing, his seat lurched forward. This could dramatically interfere with a pilot’s control of the aircraft (“he practically ended up on top of the controls... This is potentially quite nasty,” P4-4). Engineers had investigated and found no defects or blockage, but that the seat rails appeared very slightly worn. These were replaced, and the issue referred to the airframe manufacturer. Investigators were waiting on a recommendation from them, so for the moment, further action on their part was considered unlikely to make any further contribution.</p>

Implications: Organizing culture, learning and expertise

Risks were responded to in IRIS by initiating and shaping distributed processes of investigation and action, and the need for this action was evaluated in terms of what it might achieve. Primarily, these findings demonstrate that IRIS not only provided a mechanism to monitor and analyze organizational risks. They equally provided an important means by which organizational culture and practice was constituted and fashioned.

In practice, incidents were actively used for organizational purposes. They were used as points around which diverse operational specialists were connected and brought

together to manage risks. They facilitated communication about risks that otherwise may not have taken place. And they provided small, concrete cues that acted as focal points for line personnel, operational specialists and managers to reflect on, and enquire into, the adequacy of operational processes—and their assumptions about them. In a real sense, incident reports were used to direct and ground the ongoing reflection and learning that seems critical to maintaining effective and informed cultures of safety (Reason, 1997; Pidgeon and O’Leary, 1994). As such, IRIS directly address some of the most insidious organizational risks that allow the causes of disaster to incubate and remain latent (Turner, 1994; Turner and Pidgeon, 1997; Reason, 1997; Vaughan, 1996): erroneous beliefs and false assumptions.

The analysis also points to how a small group of investigators can lead broader processes of organizational learning. Investigators chose many of the issues that were explored by operational specialists and they shaped the nature of the ensuing investigations. They linked their processes of interpreting incidents with the reflection and learning of personnel around the organization. In this sense, processes of risk management involved the bridging of individual and organizational learning (e.g. Holmqvist, 2003).

Further, assessments of risk were inextricably tied up with judgements of how those risks could be acted on and resolved—and by whom. The practical interplay between processes of risk assessment and risk evaluation has received limited empirical scrutiny to date (Wilpert and Fahlbruch, 2002; Hutter and Lloyd-Bostock, 1990). Here, interpreting risk was heavily oriented to acting, and was structured by a set of criteria that evaluated the level of achievable action and the likely value of the results. While in many domains, risks have traditionally been evaluated in purely negative terms—stopping the unacceptable and preventing the intolerable—in practice here, risks were evaluated as sources of potential gains and valuable improvements to organizational processes. One important implication of this is how judgements of risk were used in practice as an interpretive tool. Risk assessments were constructs, constructed for the purpose of directing attention, focusing resources and getting action. Like Smith’s (1988, 1989) definition of managerial ‘problems’, assessments of risk were applied in order to generate and direct organizational action (e.g. Power, 2003).

Practically, these findings reaffirm the importance of social skills in risk management (e.g. Weick and Roberts, 1993; Weick, Sutcliffe and Obstfeld, 1999). For instance, the practical ability to frame and pose appropriate questions to initiate and shape the enquiry of others seems a key, if easily overlooked, skill. IRIS were used to create an organizational space in which this collective reflection and participative learning could legitimately unfold. Further examination of how participative engagement in risk management unfolds (Hutter, 2001), and how these social processes are supported, would seem a valuable area for future research. For instance, one important issue is the degree to which formal risk management systems can either support or undermine worker participation. This research echoes past work that suggests both the formal design and informal culture of safety oversight are crucial in this regard (Hopkins, 2005; Vaughan, 1990; Hutter, 2001). Participation in the investigation and management of risks requires issues of blame and trust to be sensitively negotiated. The fear of blame for risks can severely impact the nature and effectiveness of risk management (Power, 2007; Jeffcott, Pidgeon and Weyman, 2006). Equally, participation in risk management requires work and therefore

encouragement. It likely depends on effectively communicating and rewarding the resulting improvements to organizational safety (Reason, 1997).

Broader applicability and limitations

Adopting an in-depth qualitative, inductive approach to examine risk management in IRIS allowed a detailed picture of practice to be developed, and a range of theoretical insights to be elaborated on. This research nonetheless has limitations, most notably regarding the potential applicability of these findings to other settings and domains.

The theoretical account of practice developed here, in terms of three core concepts, represent tentative contributions towards the development of a more advanced theoretical understanding of this area (Weick, 1995b). Equally, it should be emphasized that practice was characterized through these three concepts to provide a degree of simplicity and clarity in capturing important facets and ideals of practice. No claims are made regarding whether such clarity or ideals have been—or perhaps ever can be—actually achieved in practice. These concepts were developed in order to provide theoretical insight, not as an assertion that the practitioners studied here in some way attained, for instance, a state of perfect and perpetual ‘interpretive vigilance’.

In terms of applicability, the findings from this research are necessarily based on the examination of a relatively focused ‘sample’ of risk management practice—that within IRIS in a small number of aviation organizations. Research that focuses on depth and detail of understanding necessarily has to sacrifice breadth (Weick, 1979), and this research did not aim to address how widespread or generalizable these findings may be. There are, nonetheless, strong grounds for the wider applicability or ‘transferability’ of the findings (King, 2000).

First, the structure of IRIS are reasonably standard across a wide range of domains (Reason, 1997; Barach and Small, 2000). It is reasonable to expect that the challenges faced and the work conducted within them are similar too. The findings may therefore provide useful explanations of these—though that, of course, is an empirical question.

Second, by focusing on questions concerning the salient features of risk management practice, and the ways in which organizational events were interpreted and responded to, the findings contribute to a more general level of understanding of the interpretive practices and organizational sensemaking involved in risk management (e.g. Weick, 1995a). These theoretical insights into the practical nature of organizational risk management may be productively transferred to other risk management settings and help explain practice there. That transferability is, again, an empirical question.

Conclusions

Studying risk management in IRIS demonstrates how distinct and complex practices of risk management can emerge in organizations. An increasing number of organizations use IRIS as a way of accessing risk data that would not otherwise be easily accessible—data arising from the minor mishaps and errors that are a normal part of organizational life. The beliefs, assumptions and tactics that supported

practices of risk management here suggest alternatives to the widely instituted predictive methods of risk analysis that focus on calculating and ranking risks according to the probability of future adverse impacts. Instead, practice here was found to be based on attempts to diagnose and respond to deficiencies in current organizational capacities for resilience.

In practice, risk was defined and handled in ways that either extended or were entirely at odds with current normative designs for risk management. Risk was both about the limits of knowledge and control in organizations, and risk management was a process focused on identifying these limits and deploying resources to address them. This extends our view of organizational risk management by broadening our definition of risk. In formal models and guidelines, risk is equated directly with future adverse consequences. In this study of practice, averting future adverse outcomes defined the aim and ends of risk management, but not the means. Risks were identified, assessed and managed in terms of a range of broader criteria and considerations. Risk assessment was focused on the capacity for underlying organizational processes to provide resilience and control in the face of error, by considering the quality and performance of sociotechnical work practices. Risk identification was concerned with flagging up areas of uncertainty and ambiguity in current models of the organization, and in current work practices. Risk resolution was effected by organizing participative investigation, action and change around these uncertainties in models and practices. As such, the main concerns were to generate new knowledge, revise old knowledge and reshape practice in light of this. In combination, these practices demonstrate risk management to be a more interactive and active process of examining the minutiae of operational practice, actively reflecting on current knowledge, engaging with specialists and experts, and facilitating organization-wide processes of enquiry, reflection, learning and change.

This paper therefore presents a first step towards characterizing the practices that underlie the core risk management tasks of risk identification, risk assessment and risk resolution in organizations. Lifting the lid on the work of risk management in this way reveals the complex, multifaceted, subtle and sometimes surprising practices that can support this increasingly important organizational function.

References

- Barach, P., and Small, D. (2000). 'Reporting and Preventing Medical Mishaps: Lessons From Non-Medical Near Miss Reporting Systems'. *British Medical Journal*, 320, 759-763.
- Cabinet Office. (2002). *Risk: Improving Government's Capability to Handle Risk and Uncertainty*. London: Cabinet Office Strategy Unit.
- Committee of the Sponsoring Organizations of the Treadway Commission. (2004). *Enterprise Risk Management—Integrated Framework*. Available online at www.coso.org/publications.htm
- Cook, S. D. N., and Brown, J. S. (1999). 'Bridging Epistemologies: The Generative Dance Between Organizational Knowledge and Organizational Knowing'. *Organization Science*, 10(4), 381-400.
- Cunha, M. P., Clegg, S.R., and Kamoche, K. (2006). 'Surprises in Management and Organization: Concept, Sources and a Typology'. *British Journal of Management*, 17, 317-329.
- Emerson, R. M., Fretz, R. I., and Shaw, L. L. (1995). *Writing Ethnographic Fieldnotes*. London: University of Chicago Press.
- Freudenburg, W. R. (1992). 'Nothing Recedes like Success? Risk Analysis and the Organizational Amplification of Risks'. *Risk: Issues in Health and Safety*, 3(1), 1-35.
- Glaser, B. G., and Strauss, A. L. (1967). *The Discovery of Grounded Theory*. Chicago: Aldine.
- Hollnagel, E. (2004). *Barriers and Accident Prevention: Or How to Improve Safety by Understanding the Nature of Accidents*. Aldershot: Ashgate.
- Hollnagel, E., Woods, D. D. and Leveson, N. (2006). *Resilience Engineering: Concepts and Precepts*. Aldershot: Ashgate.
- Holmqvist, M. (2003). 'A Dynamic Model of Intra- and Interorganizational Learning'. *Organization Studies*, 24(1), 95-123.
- Hood, C., and Jones, D.K.C. (Eds.). (1996). *Accident and Design: Contemporary Debates on Risk Management*. London: Taylor and Francis.
- Hopkins, A. (2005). *Safety, Culture and Risk: The Organizational Causes of Disasters*. Sydney: CCH.
- Hutter, B. (2001). *Regulation and Risk: Occupational Health and Safety on the Railways*. Oxford: Oxford University Press.
- Hutter, B., and Lloyd-Bostock, S. (1990). 'The Power of Accidents: The Social and Psychological Impact of Accidents and the Enforcement of Safety Regulations'. *The British Journal of Criminology*, 30(4), 409-422.

Institute of Risk Management. (2002). *A Risk Management Standard*. London: IRM, ALARM and AIRMIC.

Jeffcott S., Pidgeon N., Weyman A. and Walls J. (2006) 'Risk, Trust and Safety Culture in UK Train Operating Companies.' *Risk Analysis*. 26(5): 1105-1121.

Kiesler, S., and Sproull, L. (1982). 'Managerial Response to Changing Environments: Perspectives on Problem Sensing from Social Cognition'. *Administrative Science Quarterly*, 27, 548-570.

King, N. (2000). 'Making Ourselves Heard: The Challenges Facing Advocates of Qualitative Research in Work and Organizational Psychology'. *European Journal of Work and Organizational Psychology*, 9(4), 589-596.

Lave, J., and Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge: Cambridge University Press.

Lee, T. W. (1999). *Using Qualitative Methods in Organizational Research*. London: Sage.

Locke, K. (2001). *Grounded Theory in Management Research*. London: Sage.

Macrae, C. (2007). *Interrogating the Unknown: Risk Analysis and Sensemaking in Airline Safety Oversight*. ESRC Centre for Analysis of Risk and Regulation Discussion Paper 43. London: London School of Economics and Political Science.

Muermann, A., and Oktem, U. (2002). 'The Near-Miss Management of Operational Risk'. *Journal of Risk Finance*, Fall, 25-36.

O'Leary, M., and Chappell, S. L. (1996). 'Confidential Incident Reporting Systems Create Vital Awareness of Safety Problems'. *International Civil Aviation Organization Journal*, 51, 11-13.

Orr, J. (1996). *Talking about Machines: An Ethnography of a Modern Job*. London: Cornell University Press.

Phimister, J. R., Oktem, U., Kleindorfer, P. R., & Kunreuther, H. (2003). 'Near-Miss Incident Management in the Chemical Process Industry'. *Risk Analysis*, 23(3), 445-459.

Pidgeon, N. (1998). 'Safety Culture: Key Theoretical Issues'. *Work and Stress*, 12(3), 202-216.

Pidgeon, N., and Henwood, K. (2003). 'Grounded Theory'. In M. Hardy, and Bryman, A. (Eds.), *Handbook of Data Analysis*. London: Sage.

Pidgeon, N., and O'Leary, M. (1994). 'Organizational Safety Culture: Implications for Aviation Practice'. In N. Johnston, McDonald, N. and Fuller, R. (Eds.), *Aviation Psychology in Practice*. Aldershot: Ashgate.

Pidgeon, N., and O'Leary, M. (2000). 'Man-Made Disasters: Why Technology and Organizations (Sometimes) Fail'. *Safety Science*, 34, 15-30.

Power, M. (2003). *The Invention of Operational Risk*. ESRC Centre for Analysis of Risk and Regulation Discussion Paper 16, London School of Economics and Political Science, London.

Power, M. (2004). *The Risk Management of Everything: Rethinking the Politics of Uncertainty*. London: Demos.

Power, M. (2007). *Organized Uncertainty: Designing a World of Risk Management*. Oxford: Oxford University Press.

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

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

Reason, J., and Hobbs, A. (2003). *Managing Maintenance Error: A Practical Guide*. Aldershot: Ashgate.

Robson, C. (2002). *Real World Research: A Resource for Social Scientists and Practitioner-Researchers*. Oxford: Blackwell Publishers.

Rochlin, G. I. (1989). 'Informal organizational networking as a crisis-avoidance strategy: US Naval flight operations as a case study'. *Industrial Crisis Quarterly*, 3, 159-176.

Sheaffer, Z., Richardson, B. and Rosenblatt, Z. (1998). 'Early-Warning-Signals Management: A Lesson from the Barings Crisis'. *Journal of Contingencies and Crisis Management*, 6(1), 1-22.

Sheffi, Y. (2005). *The Resilience Enterprise: Overcoming Vulnerability for Competitive Advantage*. Cambridge: MIT Press.

Smith, G. F. (1988). 'Towards a Heuristic Theory of Problem Structuring'. *Management Science*, 34(12), 1489-1506.

Smith, G. F. (1989). 'Defining Managerial Problems: A Framework for Prescriptive Theorizing'. *Management Science*, 35(8), 963-980.

Smithson, M. (1989). *Ignorance and Uncertainty: Emerging Paradigms*. London: Springer-Verlag.

Sutcliffe, K. M., and Vogus, T. J. (2003). 'Organizing for Resilience'. In K. S. Cameron, and Dutton, J. E. (Eds.), *Positive Organizational Scholarship: Foundations of a New Discipline*. London: Berrett-Koehler.

Svenson, O. (1991). 'The Accident Evolution and Barrier Function (AEB) Model Applied to Incident Analysis in the Processing Industries'. *Risk Analysis*, 11, 499-507.

Turner, B. (1978). *Man-Made Disasters*. London: Wykeham.

- Turner, B. (1983). 'The Use of Grounded Theory for the Qualitative Analysis of Organizational Behaviour'. *Journal of Management Studies*, 20(3), 333-348.
- Turner, B. (1994). 'Causes of Disaster: Sloppy Management'. *British Journal of Management*, 5, 215-219.
- Turner, B., and Pidgeon, N. (1997). *Man-Made Disasters* (2nd ed.). Oxford: Butterworth-Heinemann.
- van der Schaaf, T. V., Lucas, D. A. and Hale, A. R. (Eds.). (1991). *Near Miss Reporting as a Safety Tool*. Oxford: Butterworth-Heinemann.
- Vaughan, D. (1990). 'Autonomy, interdependence, and social control: NASA and the space shuttle Challenger'. *Administrative Science Quarterly*, 35, 225-257.
- Vaughan, D. (1996). *The Challenger Launch Decision: Risky Technology, Culture and Deviance at NASA*. London: Chicago University Press.
- Weick, K. E. (1979). *The Social Psychology of Organizing* (2nd ed.). London: McGraw-Hill.
- Weick, K. E. (1987). 'Organizational Culture as a Source of High Reliability'. *California Management Review*, 29(2), 112-127.
- Weick, K. E. (1995a). *Sensemaking in Organizations*. London: Sage.
- Weick, K. E. (1995b). 'What Theory is Not, Theorizing is'. *Administrative Science Quarterly*, 40, 385-390.
- Weick, K. E., and Roberts, K. H. (1993). 'Collective Mind in Organizations: Heedful Interrelating on Flight Decks'. *Administrative Science Quarterly*, 38(357-381).
- Weick, K. E., and Sutcliffe, K. M. (2001). *Managing the Unexpected: Assuring High Performance in an Age of Complexity*. San Francisco: Jossey Bass.
- Weick, K. E., Sutcliffe, K. M. and Obstfeld, D. (1999). 'Organizing for High Reliability: Processes of Collective Mindfulness'. *Organizational Behaviour*, 21, 81-123.
- Wright, D. A., and Lyons, J. E. (2001). *Flight data monitoring: Its place within a safety management system*. Paper presented at the 13th European aviation safety seminar: Toward a safer Europe, Amsterdam, Netherlands.
- Wildavsky, A. (1988). *Searching for Safety*. Oxford: Transaction Publishers.
- Wilpert, B., and Fahlbruch, B. (Eds.). (2002). *System Safety: Challenges and Pitfalls of Intervention*. London: Pergamon.