Injury Control & Safety Promotion
1566-0974/00/US$ 15.00

Injury Control & Safety Promotion – 2000, Vol. 7, No. 1, pp. 3-14

Accepted 15 November 1999

Safety paradoxes and safety culture

James Reason

Department of Psychology, University of Manchester, U.K.

Abstract This paper deals with four safety paradoxes: (1) Safety is defined and measured more by its absence than its presence. (2) Defences, barriers and safeguards not only protect a system, they can also cause its catastrophic breakdown. (3) Many organisations seek to limit the variability of human action, primarily to minimise error, but it is this same variability – in the form of timely adjustments to unexpected events – that maintains safety in a dynamic and changing world. (4) An unquestioning belief in the attainability of absolute safety can seriously impede the achievement of realisable safety goals, while a preoccupation with failure can lead to high reliability. Drawing extensively upon the study of high reliability organisations (HROs), the paper argues that a collective understanding of these paradoxes is essential for those organisations seeking to achieve an optimal safety culture. It concludes with a consideration of some practical implications.

Key words Safety promotion; culture; defences; errors; adaptability; beliefs; psychological factors; human behaviour

Correspondence and reprint requests to:
James Reason
Dept. Psychology
Univ. Manchester
Oxford Road
Manchester M13 9PL
England, U.K.
Tel.: +44 161 275 2551
Fax: +44 161 275 2622
E-mail: reason@fs4.psy.man.ac.uk

Introduction A paradox is *'a statement contrary to received opinion; seemingly absurd though perhaps well-founded'* (Concise Oxford Dictionary). This paper contends that the pursuit of safety abounds with paradox, and that this is especially true of efforts to achieve a safer organisational culture. In safety, as in other highly interactive spheres, things are not always what they seem. Not only can they be contrary to surface appearances, they can also run counter to some of our most cherished beliefs. The better we understand these paradoxes, the more likely we are to create and sustain a truly safe culture.

A safe culture is an informed culture, one that knows continually where the 'edge' is without necessarily having to fall over it. The 'edge' lies between relative safety and unacceptable danger. In many industries, proximity to the 'edge' is the zone of greatest peril and also of greatest profit.[1] Navigating this area requires considerable skill on the part of system managers and operators. Since such individuals come and go, however, only a safe culture can provide any degree of lasting protection.

Simply identifying the existence of a paradox is not enough. Unlike the 'pure' sciences, in which theories are assessed by how much empirical activity they provoke, the insights of safety scientists and safety practitioners are ultimately judged by the extent to which their practical application leads to safer systems. Each of the paradoxes considered below has important practical implications for the achievement of a safe culture. Indeed, it will be argued that a shared understanding of these paradoxes is a prerequisite for acquiring an optimal safety culture.

Most of the apparent contradictions discussed in this paper have been revealed not so much by the investigation of adverse events – a topic that comprises the greater part of safety research – as from the close observation of high reliability organisations (HROs). Safety has both a negative and a positive face. The former is revealed by accidents with bad outcomes. Fatalities, injuries and environmental damage are conspicuous and readily quantifiable occurrences. Avoiding them as far as possible is the objective of the safety sciences. It is hardly surprising, therefore, that this darker face has occupied so much of our attention and shaped so many of our beliefs about safety. The positive face, on the other hand, is far more secretive. It relates to a system's intrinsic resistance to its operational hazards. Just as medicine knows more about pathology than health, so also do the safety sciences understand far more about how bad events happen than about how human actions and organisational processes also lead to their avoidance, detection and containment. It is this imbalance that has largely created the paradoxes.

The remainder of the paper is in six parts. The next section previews the four safety paradoxes to be considered here. The ensuing four sections each consider one of these safety paradoxes in more detail. The concluding section summarises the practical implications of these paradoxes for achieving and preserving a safer culture.

Previewing the safety paradoxes

- Safety is defined and measured more by its absence than by its presence.
- Measures designed to enhance a system's safety – defences, barriers and safeguards – can also bring about its destruction.
- Many, if not most, engineering-based organisations believe that safety is best achieved through a predetermined consistency of their processes and behaviours, but it is the uniquely human ability to vary and adapt actions to suit local conditions that preserves system safety in a dynamic and uncertain world.
- An unquestioning belief in the attainability of absolute safety (zero accidents or target zero) can seriously impede the achievement of realisable safety goals.

A further paradox embodies elements from all of the above. If an organisation is convinced that it has achieved a safe culture, it almost certainly has not. Safety culture, like a state of grace, is a product of continual striving. There are no final victories in the struggle for safety.

The first paradox: how safety is defined and assessed

The Concise Oxford Dictionary defines safety as '*freedom from danger and risks*'. But this tells us more about what comprises 'unsafety' than about the substantive properties of safety itself. Such a definition is clearly unsatisfactory. Even in the short term, as during a working day or on a particular journey, we can never escape danger – though we may not experience its adverse consequences in that instance. In the longer term, of course, most of the risks and hazards that beset human activities are universal constants. Gravity, terrain, weather, fire and the potential for uncontrolled releases of mass, energy and noxious substances are ever-present dangers. So, in the strict sense of the definition, we can never be safe. A more appropriate definition of safety would be 'the ability of individuals or organisations to deal with risks and hazards so as to avoid damage or losses and yet still achieve their goals'.

Even more problematic, however, is that safety is measured by its occasional absences. An organisation's safety is commonly assessed by the number and severity of negative outcomes (normalised for exposure) that it experiences over a given period. But this is a flawed metric for the reasons set out below.

First, the relationship between intrinsic 'safety health' and negative outcomes is, at best, a tenuous one. Chance plays a large part in causing bad events – particularly so in the case of complex, well-defended technologies.[2] As long as hazards, defensive weaknesses and human fallibility continue to co-exist, unhappy chance can combine them in various ways to bring about a bad event. That is the essence of the term 'accident'. Even the most resistant organisations can suffer a bad accident. By the same token, even the most vulnerable systems can evade disaster, at least for a time. Chance does not take sides. It afflicts the deserving and preserves the unworthy.

Second, a general pattern in organisational responses to a safety management programme is that negative outcome data decline rapidly at first and then gradually bottom out to some asymptotic value. In commercial aviation, for example, a highly safety conscious industry, the fatal accident rate has remained relatively unchanged for the past 25 years.[3] Comparable patterns are found in many other domains. During the period of rapid decline, it seems reasonable to suppose that the marked diminution in accident rates actually does reflect some improvement in a system's intrinsic 'safety health'. But once the plateau has been reached, periodic variations in accident rates contain more noise than valid safety signals. At this stage of an organisation's safety development, negative outcome data are a poor indication of its ability to withstand adverse events in the future. This is especially true of well-defended systems such as commercial aviation and nuclear power generation that are, to a large extent, victims of their own success. By reducing accident rates to a very low level they have largely run out of 'navigational aids' by which to steer towards some safer state.

The diminution in accident rates that is apparent in most domains is a product not only of local safety management efforts, but also of a growing public intolerance for third-party risks, environmental damage and work-related injuries. This, in turn, has led to increasingly comprehensive safety legislation in most industrialised nations. Even in the least responsible organisations, merely keeping one step ahead of the regulator requires the implementation of basic safety measures that are often sufficient to bring about dramatic early reductions in accident rates. The important issue, however, is what happens once the plateau has been reached. It is at this point that an organisation's safety culture takes on a profound significance. Getting from bad to average is relatively easy; getting from average to excellent is very hard. And it is for the latter purpose that an understanding of the paradoxes is crucial.

In summary: while high accident rates may reasonably be taken as indicative of a bad safety state, low asymptotic rates do not necessarily signal a good one. This asymmetry in the meaning of negative outcome data lies at the heart of many of the subsequent paradoxes to be discussed later. It also has far-reaching cultural implications. There are at least two ways to interpret very low or nil accident rates in a given accounting period. A very common one is to believe that the organisation actually has achieved a safe state: that is, it takes no news as good news and sends out congratulatory messages to its workforce. High-reliability organisations, on the other hand, become worried, accepting that no news really is no news, and so adopt an attitude of increased vigilance and heightened defensiveness.[4,5]

The second paradox: dangerous defences A theme that recurs repeatedly in accident reports is that measures designed to enhance a system's safety can also bring about its destruction. Since this paradox has been discussed at length elsewhere,[6,7] we will focus on its cultural implications. Let us start with some examples of defensive failures that cover a range of domains.

- The Chernobyl disaster had its local origins in an attempt to test an electrical safety device designed to overcome the interruption of power to the emergency core cooling system that would ensue immediately after the loss of off-site electricity and before the on-site auxiliary generators were fully operative.[8]
- The advanced automation present in many modern technologies was designed, in part, to eliminate opportunities for human error. Experience in several domains, however, has shown that automation can create mode confusions and decision errors that can be more dangerous than the slips and lapses it was intended to avoid.[9,10]
- Emergency procedures are there to guide people to safety in the event of a dangerous occurrence. In a number of instances, however, strict compliance with safety procedures has killed people. On Piper Alpha, the North Sea gas and oil platform that exploded in 1988, most of the 165 rig workers that died complied strictly with the safety drills and assembled in the accommodation area. Tragically, this was directly in line with a subsequent explosion.[11] The few fire-fighters that survived the Mann Gulf forest fire disaster in 1949 dropped their heavy tools and ran, while those who died obeyed the organisational instruction to keep their fire-fighting tools with them at all times.[12]
- Personal protective equipment can save many lives, but it can also pose a dangerous threat to certain groups of people. Swedish traffic accident studies have revealed that both elderly female drivers and infants in backward-facing seats have been killed by rapidly inflating airbags following a collision.[13]

- Finally, perhaps the best example of the defence paradox is that maintenance activities – intended to repair and forestall technical failures – are the largest single source of human factors problems in the nuclear power industry.[14,15] In commercial aviation, quality lapses in maintenance are the second most significant cause of passenger deaths.[16]

There is no single reason why defences are so often instrumental in bringing about bad events. Errors in maintenance, for example, owe their frequency partly to the hands- on, high-opportunity nature of the task, and partly to the fact that certain aspects of maintenance, particularly installation and reassembly, are intrinsically error-provoking regardless of who is doing the job.[6] But some of the origins of the defensive paradox have strong cultural overtones. We can summarise these cultural issues under three headings: the trade-off problem, the control problem and the opacity problem.

THE TRADE-OFF PROBLEM An important manifestation of an organisation's cultural complexion is the characteristic way it resolves conflicts. Virtually all of the organisations of concern here are in the business of producing something: manufactured goods, energy, services, the extraction of raw materials, transportation and the like. All such activities involve the need to protect against operational hazards. A universal conflict, therefore, is that between production and protection. Both make demands upon limited resources. Both are essential. But their claims are rarely perceived as equal. It is production rather than protection that pays the bills, and those who run these organisations tend to possess productive rather than protective skills. Moreover, the information relating to the pursuit of productive goals is continuous, credible and compelling, while the information relating to protection is discontinuous, often unreliable, and only intermittently compelling (i.e., after a bad event). It is these factors that lie at the root of the trade-off problem. This problem can best be expressed as that of trading protective gains for productive advantage. It has also been termed risk homeostasis[17] or risk compensation – the latter term is preferable since it avoids some of Wilde's more controversial assumptions.[18]

The trade-off problem has been discussed at length elsewhere.[18-20] Just one example will be sufficient to convey its essence. The Davy lamp, invented in 1815, was designed to isolate the light source, a naked flame, from the combustible gases present in mines. But the mine owners were quick to see that it also allowed miners to work on seams previously regarded as too dangerous. The incidence of mine explosions increased dramatically, reaching a peak in the 1860s.[20] Improvements in protection afforded by technological developments are often put in place during the aftermath of a disaster. Soon, however, this increased protection is seen as offering commercial advantage, leaving the organisation with the same or even less protection than it had previously.

THE CONTROL PROBLEM Another challenge facing all organisations is how to restrict the enormous variability of human behaviour to that which is both productive and safe. Organisational managers have a variety of means at their disposal:[21,22] administrative controls (prescriptive rules and procedures), individual controls (selection, training and motivators), group controls (supervision, norms and targets) and technical controls (automation, engineered safety features, physical barriers). In most productive systems, all of these controls are used to some degree; but the balance between them is very much a reflection of the organisational culture. What concerns us here, however, is the often disproportionate reliance placed upon prescriptive procedures.

Standard operating procedures are necessary. This is not in dispute. Since people change faster than jobs, it is essential that an organisation's collective wisdom is recorded and passed on. But procedures are not without problems, as indicated by some of the examples listed above. They are essentially feed-forward control devices – prepared at one time and place to be applied at some future time and place – and they suffer, along with all such control systems, the problem of dealing with local variations. Rule-based controls can encounter at least three kinds of situation: those in which they are correct and appropriate, those in which they are inapplicable due to local conditions, and those in which they are absent entirely. A good example of the latter is the predicament facing Captain Al Haynes and his crew in United 232 when he lost all three hydraulic systems on his DC10 due to the explosion of his tail-mounted, number two engine.[23] The probability of losing all three hydraulic systems was calculated at one in a billion, and there were no procedures to cover this unlikely emergency. Far more common, however, are situations in which the procedures are unworkable, incomprehensible or simply wrong. A survey carried out in the US nuclear industry, for example, identified poor procedures as a factor in some 60% of all human performance problems.[15]

There is a widespread belief among the managers of highly proceduralised organisations that suitable training, along with rigid compliance, should eliminate the vast majority of human unsafe acts. When such errors and violations do occur, they are often seen as moral issues warranting sanctions. But, for the most part, punishing people does not eliminate the systemic causes of their unsafe acts. Indeed, by isolating individual actions from their local context, it can impede their discovery.

THE OPACITY PROBLEM In the weeks following some foreign technological disaster, we often hear our country's spokespeople claiming that it couldn't happen here because our barriers and safeguards are so much more sophisticated and extensive. This assertion captures an important consequence of the opacity problem: the failure to realise that defences, particularly defences-in-depth, can create and conceal dangers as well as protect against them. When this ignorance leads to a collective belief in the security of high-technology systems, the problem takes on cultural significance.

Defences-in-depth are created by diversity and redundancy. Barriers and safeguards take many forms. 'Hard' defences include automated safety features, physical containment, alarms and the like. 'Soft' defences include rules and procedures, training, drills, briefings, permit-to-work systems and many other measures that rely heavily on people and paper. This assortment of safety-enhancing measures is widely distributed throughout the organisation. This makes such extensively defended systems especially vulnerable to the effects of an adverse safety culture. Only culture can reach equally into all parts of the system and exert some consistent effect, for good or ill.[24]

While such diversity has undoubtedly enhanced the security of high-technology systems, the associated redundancy has proved to be a mixed blessing. By increasing complexity, it also makes the system more opaque to those who manage and control it.[7,25,26] The opacity problem takes a variety of forms.

- Operator and maintainer failures may go unnoticed because they are caught and concealed by multiple backups.[27]
- Such concealment allows undiscovered errors and latent conditions (resident pathogens) to accumulate insidiously over time, thus increasing the possibility of inevitable weaknesses in the defensive layers lining up to permit the passage of an accident trajectory.[6,28]
- By adding complexity to the system, redundant defences also increase the likelihood of unforeseeable common-mode failures. While the assumption of independence may be appropriate for purely technical failures, errors committed by managers, operators and maintainers are uniquely capable of creating problems that can affect a number of defensive layers simultaneously. At Chernobyl, for example, the operators successively disabled a number of supposedly independent, engineered safety features in pursuit of their testing programme.

Dangerous concealment combined with the obvious technological sophistication of redundant defences can readily induce a false sense of security in system managers, maintainers and operators. In short, they forget to be afraid – or, as in the case of the Chernobyl operators, they never learn to be afraid. Such complacency lies on the opposite pole from a safe culture.

The third paradox: consistency versus variability Hollnagel[20] conducted a survey of the human factors literature to identify the degree to which human error has been implicated in accident causation over the past few decades. In the 1960s, when the problem first began to attract serious attention, the estimated contribution of human error was around 20%. By the 1990s, this figure had increased fourfold to around 80%. One of the possible reasons for this apparent growth in human fallibility is that accident investigators are now far more conscious that contributing errors are not confined to the 'sharp end' but are present at all levels of a system, and even beyond. Another is that the error causal category has, by default, moved more and more into the investigatory spotlight due to great advances in the reliability of mechanical and electronic components over the past forty years.

Whatever the reason, the reduction – or even elimination – of human error has now become one of the primary objectives of system managers. Errors and violations are viewed, reasonably enough, as deviations from some desired or appropriate behaviour. Having mainly an engineering background, such managers attribute human unreliability to unwanted variability. And, as with technical unreliability, they see the solution as one of ensuring greater consistency of human action. They do this, as we have seen, through procedures and by buying more automation. What they often fail to appreciate, however, is that human variability in the form of moment-to-moment adaptations and adjustments to changing events is also what preserves system safety in an uncertain and dynamic world. And therein lies the paradox. By striving to constrain human variability, they are also undermining one the system's most important safeguards.

The problem has been encapsulated by Weick's insightful observation[5] that 'reliability is a dynamic non-event.' It is dynamic because processes remain under control due to compensations by human components. It is a non-event because safe outcomes claim little or no attention. The paradox is rooted in the fact that accidents are salient, while non-events, by definition, are not. Almost all of our methodological tools are geared to investigating adverse events. Very few of them are suited to creating an understanding of why timely adjustments are necessary to achieve successful outcomes in an uncertain and dynamic world

Recently, Weick et al.[4] challenged the received wisdom that an organisation's reliability depends upon the consistency, repeatability and invariance of its routines and activities. Unvarying performance, they argue, cannot cope with the unexpected. To account for the success of high reliability organisations (HROs) in dealing with unanticipated events, they distinguish two aspects of organisational functioning: cognition and activity. The cognitive element relates to being alert to the possibility of unpleasant surprises and having the collective mindset necessary to detect, understand and recover them before they bring about bad consequences. Traditional 'efficient' organisations strive for stable activity patterns yet possess variable cognitions – these differing cognitions are most obvious before and after a bad event. In HROs, on the other hand, 'there is variation in activity, but there is stability in the cognitive processes that make sense of this activity'.[4] This cognitive stability depends critically upon an informed culture – or what Weick and his colleagues have called 'collective mindfulness'.

Collective mindfulness allows an organisation to cope with the unanticipated in an optimal manner. 'Optimal' does not necessarily mean 'on every occasion', but the evidence suggests that the presence of such enduring cognitive processes is a critical component of organisational resilience. Since catastrophic failures are rare events, collectively mindful organisations work hard to extract the most value from what little data they have. They actively set out to create a reporting culture by commending, even rewarding, people for reporting their errors and near misses. They work on the assumption that what might seem to be an isolated failure is likely to come from the confluence of many 'upstream' causal chains. Instead of localising failures, they generalise them. Instead of applying local repairs, they strive for system reforms. They do not take the past as a guide to the future. Aware that system failures can take a wide variety of yet-to-be-encountered forms, they are continually on the lookout for 'sneak paths' or novel ways in which active failures and latent conditions can combine to defeat or by-pass the system defences. In short, HROs are preoccupied with the possibility of failure – which brings us to the last paradox to be considered here.

The fourth paradox: target zero Some years ago, US Vice-President Al Gore declared his intention of eradicating transport accidents. Comparable sentiments are echoed by the top managers of by-the-book companies, those having what Westrum[29] has called 'calculative' cultures. They announce a corporate goal of 'zero accidents' and then set their workforce the task of achieving steadily diminishing accident targets year by year – what I have earlier termed the 'negative production' model of safety management.

It is easy to understand and to sympathise with such goal-setting. A truly committed management could hardly appear to settle for anything less. But 'target zero' also conveys a potentially dangerous misrepresentation of the nature of the struggle for safety: namely, that the 'safety war' could end in a decisive victory of the kind achieved by a Waterloo or an Appomattox. An unquestioning belief in victory can lead to defeat in the 'safety war'. The key to relative success, on the other hand, seems to be an abiding concern with failure

HROs see the 'safety war' for what it really is: an endless guerrilla conflict. They do not seek a decisive victory, merely a workable survival that will allow them to achieve their productive goals for as long as possible. They know that the hazards will not go away, and accept that entropy defeats all systems in the end. HROs accept setbacks and nasty surprises as inevitable. They expect to make errors and train their workforce to detect and recover them. They constantly rehearse for the imagined scenarios of failure and then go on to brainstorm novel ones. In short, they anticipate the worst and equip themselves to cope with it.

A common response to these defining features of HROs is that they seem excessively bleak. 'Doom-laden' is a term often applied to them. Viewed from a personal perspective, this is an understandable reaction. It is very hard for any single individual to remain ever mindful of the possibility of failure, especially when such occurrences have personal significance only on rare occasions. No organisation is just in the business of being safe. The continuing press of productive demands is far more likely to engage the forefront of people's minds than the possibility of some unlikely combination of protective failures. This is exactly why safety culture is so important. Culture transcends the psychology of any single person. Individuals can easily forget to be afraid. A safe culture, however, can compensate for this by providing the reminders and ways of working that go to create and sustain intelligent wariness. The individual burden of chronic unease is also made more supportable by knowing that the collective concern is not so much with the occasional – and inevitable – unreliability of its human parts, as with the continuing resilience of the system as a whole.

The practical implications By what means can we set about transforming an average safety culture into an excellent one? The answer, I believe, lies in recognising that a safe culture is the product of a number of inter-dependent sub-cultures, each of which – to some degree – can be socially engineered. An *informed culture* can only be built on the foundations of a *reporting culture*. And this, in turn, depends upon establishing a *just culture*. In this concluding section, we will look at how to build these two sub-cultures. The other elements of a safe culture – a *flexible culture* and a *learning culture* – hinge largely upon the establishment of the previous two. They have been discussed at length elsewhere[5,6] and will not be considered further here.

In the absence of frequent bad outcomes, knowledge of where the 'edge' lies can only come from persuading those at the human-system interface to report their 'free lessons'. These are the mostly inconsequential errors, incidents and near misses that could have caused injury or damage. But people do not readily confess their blunders, particularly if they believe such reports could lead to disciplinary action. Establishing trust, therefore, is the first step in engineering a reporting culture – and this can be very big step. Other essential characteristics are that the organisation should possess the necessary skills and resources to collect, analyse and disseminate safety-related information and, crucially, it should also have a management that is willing to act upon and learn from these data.

A number of effective reporting systems have been established, particularly in aviation. Two behavioural scientists involved in the creation of two very successful systems, the Aviation Safety Reporting System developed by NASA and the British Airways Safety Information System, have recently collaborated to produce a blueprint for engineering a reporting culture.[30] The main features are summarised below.

- A qualified indemnity against sanctions – though not blanket immunity.
- A reliance on confidentiality and de-identification rather than complete anonymity.
- The organisational separation of those who collect and analyse the data from those responsible for administering sanctions.
- Rapid, useful and intelligible feedback – after the threat of punishment, nothing deters reporters more than a lack of any response.
- Reports should be easy to make. Free text accounts appear to be more acceptable to reporters than forced-choice questionnaires.

The first three of these measures relate to the issue of punishment. In the past, many organisations relied heavily upon the threat of sanctions to shape reliable human behaviour. More recently, the pendulum has swung towards the establishment of 'no blame' cultures. But like the excessively punitive culture it supplanted, this approach is neither desirable nor workable. A small proportion of unsafe acts are indeed reckless and warrant severe sanctions. What is needed is a just culture, one in which everyone knows where the line must be drawn between acceptable and unacceptable actions. When this is done, the evidence suggests that only around 10% of unsafe acts fall into the unacceptable category.[6,31] This means that around 90% of unsafe acts are largely blameless and could be reported without fear of punishment.

So how should this line be drawn? Many organisations place the boundary between errors and procedural violations, arguing that only the latter are deliberate actions. But there are two problems with this: some errors arise from unacceptable behaviours, while some violations are enforced by organisational rather than by individual shortcomings, and so should not be judged as unacceptable. Marx[31] has proposed a better distinction. The key determinant of blameworthiness, he argues, is not so much the act itself – error or violation – as the nature of the behaviour in which it was embedded. Did this behaviour involve unwarranted risk-taking? If so, then the act would be blameworthy regardless of whether it was an error or a violation. Often, of course, the two acts are combined. For instance, a person may violate procedures by taking on a double shift and make a dangerous mistake in the final hour. Such an individual would merit punishment because he or she took an unjustifiable risk in working a continuous 18 hours, thus increasing the likelihood of an error.[32]

These are fine judgements and there is insufficient space to pursue them further here. The important point, however, is that such determinations – ideally involving both management and peers – lie at the heart of a just culture. Without a shared agreement as to where such a line should be drawn, there can never be an adequate reporting culture. Without a reporting culture, there could not be an informed culture. It is the knowledge so provided that gives an optimal safety culture its defining characteristics: a continuing respect for its operational hazards, the will to combat hazards in a variety of ways and a commitment to achieving organisational resilience. And these, I have argued, require a 'collective mindfulness' of the paradoxes of safety.

References

1 Hudson PTW. Psychology and safety. Leiden: University of Leiden, 1997.

2 Reason J. Achieving a safe culture: theory and practice. Work & Stress 1998;12:293-306.

3 Howard RW. Breaking through the 10^6 barrier. Proc Int Fed Airworthiness Conf, Auckland, NZ, 20-23 October 1991.

4 Weick KE, Sutcliffe KM, Obstfeld D. Organizing for high reliability: processes of collective mindfulness. In: Staw B, Sutton R, editors. Research in Organizational Behavior 1999;21:23-81.

5 Weick KE. Organizational culture as a source of high reliability. Calif Management Rev 1987;29:112-27.

6 Reason J. Managing the risks of organizational accidents. Aldershot: Ashgate, 1997.

7 Rijpma JA. Complexity, tight-coupling and reliability: Connecting normal accidents theory and high reliability theory. J Conting Crisis Management 1997;5:15-23.

8 Medvedev G. The truth about Chernobyl. New York: Basic Books, 1991.

9 Sarter NB, Woods DD. Mode error in the supervisory control of automated systems. Proc Human Factors Soc 36th Annual Meeting. Atlanta, October 1992.

10 Hughes D. Incidents reveal mode confusion. Aviation Week & Space Technol, 30 January 1995: 5.

11 Punchard E. Piper Alpha: a survivor's story. London: W.H.Allen, 1989.

12 Weick KE. The collapse of sensemaking in organizations. Admin Sci Q 1993;38:628-52.

13 Farquhar B. Safety among different subpopulations. Proc Eur Conf Safety in the Modern Society, Helsinki, Finland, 14-15 September 1999.

14 Rasmussen J. What can be learned from human error reports? In: Duncan K, Gruneberg M, Wallis D, editors.

Changes in working life. London: Wiley, 1980.
15 INPO. An analysis of root causes in 1983 and 1984 significant event reports. Atlanta: Inst Nuclear Power Operations, 1985.
16 Davis RA. Human factors in the global market place. Proc Annual Meeting Human Factors & Ergon Soc. Seattle, 12 October 1993.
17 Wilde GJS. The theory of risk homeostasis: implications for safety and health. Risk Anal 1982;2:209-55.
18 Evans L. Traffic safety and the driver. New York: Van Nostrand, 1991.
19 Adams J. Risk and freedom. London: Transport Publishing Projects, 1985.
20 Hollnagel E. Human reliability analysis: context and control. London: Acad Press, 1993.
21 Hopwood AG. Accounting systems and managerial behaviour. Hampshire, UK: Saxon House, 1974.
22 Reason J, Parker D, Lawton R. Organizational controls and safety: the varieties of rule-related behaviour. J Occup Organizational Psychol 1998; 71:289-304.
23 Haynes AC. United 232: coping with the loss of all flight controls. Flight Deck 1992;3:5-21.
24 Reason J. Human error. New York: Cambridge Univ Press, 1990.
25 Perrow C. Normal accidents: Living with high-risk technologies. New York: Basic Books, 1984.
26 Sagan SD. The limits of safety: organizations, accidents and nuclear weapons. Princeton, NJ: Princeton Univ Press, 1994.
27 Rasmussen J. Learning from experience? Some research issues in industrial risk management. In: Wilpert B, Qvale T, editors. Reliability and safety in hazardous work systems. Hove: LEA, 1993.
28 Turner B. Man-made disasters. London: Wykeham Publ, 1978.
29 Westrum R. Cultures with requisite imagination. In: Wise J, Hopkins V, Stager P, editors. Verification and Validation of Complex Systems. Berlin: Springer-Verlag, 1993.
30 O'Leary M, Chappell SL. Confidential incident reporting systems create vital awareness of safety problems. ICAO J 1996;51:11-13.
31 Marx D. Discipline: the role of rule violations. Ground Effects 1997;2:1-4.
32 Marx D. Maintenance error causation. Washington: FAA Office Aviation Med, 1999.