

Reasons for the Failure of Incident Reporting in the Healthcare and Rail Industries

C.W. Johnson

Department of Computing Science, University of Glasgow,
Glasgow, Scotland, G12 9QQ.
johnson@dcs.gla.ac.uk
<http://www.dcs.gla.ac.uk/~johnson>

Abstract. Incident reporting systems have recently been established across the UK rail and healthcare industries. These initiatives have built on the perceived success of reporting systems within aviation. There is, however, a danger that the proponents of these schemes have significantly over-estimated the impact that they can have upon the operation of complex, safety-critical systems. This paper, therefore, provides a brief overview of the problems that limit the utility of incident reporting in the the rail and healthcare industries.

1 Introduction

On the 30th November 1999, the Deputy Prime Minister, John Prescott, announced that the Confidential Incident Reporting and Analysis System (CIRAS) would be extended from the Scottish railway system to cover the entire network [4]. On the 13th June 2000, the Health Secretary, Alan Milburn, and the Chief Medical Officer for England, Liam Donaldson, announced the establishment of a centralised reporting facility for adverse incidents across the UK National Health Service (NHS) [5]. The Chief Medical Officer said; “At the moment there is no way of knowing whether the lessons learned from an incident in one part of the NHS are properly shared with the whole health service”. The Health Secretary said; “Patients, staff and the public have the right to expect the NHS to learn from its mistakes so we can ensure the alarm bells ring when there are genuine concerns so they can be nipped in the bud”. The more detailed statements that followed these media announcements made a number of claims about the benefits of incident reporting. These can be summarised as follows:

1. Incident reports help to find out why accidents do not occur. Many incident reporting forms identify the barriers that prevent adverse situations from developing into a major accident. These insights are very important. They help analysts to identify where additional support is required in order to guarantee the future benefits of those safeguards.
2. The higher frequency of incidents permits quantitative analysis. It can be argued that many accidents stem from atypical situations. They, therefore, provide relatively little information about the nature of future failures. In

contrast, the higher frequency of incidents provides greater insights into the relative proportions of particular classes of human ‘error’, systems failure, regulatory weakness etc.

3. They provide a reminder of hazards. Incident reports provide a means of monitoring potential problems as they recur during the lifetime of an application. The documentation of these problems increases the likelihood that recurrent failures will be noticed and acted upon.
4. Feedback keeps staff ‘in the loop’. Incident reporting schemes provide a means of encouraging staff participation in safety improvement. In a well-run system, they can see that their concerns are treated seriously and are acted upon by the organisation. Greater insight into national and global safety issues can be gained.
5. Data (and lessons) can be shared. Incident reporting systems provide the raw data for comparisons both within and between industries. If common causes of incidents can be observed then, it is argued, common solutions can be found. However, in practice, the lack of national and international standards for incident reporting prevents designers and managers from gaining a clear view of the relative priorities of such safety improvements.
6. Incident reporting schemes are cheaper than the costs of an accident. There is an argument that the relatively low costs of managing an incident reporting scheme should be offset against the costs of failing to prevent an accident. This is a persuasive argument. However, there is also a concern that punitive damages may be levied if an organisation fails to act upon the causes of an incident that subsequently contribute towards an accident.

Numerous studies have described tools and techniques that are intended to help realise these potential benefits of incident reporting systems [33, 51]. Very few papers analyse the reasons why some initiatives have failed to yield safety improvements [27]. This is a significant omission. Some reporting systems only elicit a very small number of contributions. Those submissions that are obtained may only come from particular sections of the workforce [7]. Internal pressures prevent safety managers from responding effectively to particular contributions [30]. The following pages, therefore, summarise the reasons for the failure of incident reporting. The analysis is illustrated by case studies that are predominantly drawn from the healthcare and rail industries. This choice is justified by recent public attention and by the different management structures that characterise these two domains. It is important to note, however, that many of these problems affect a range of other industries, including aviation [27].

2 Unrealistic Expectations

It can often be a surprise to learn of the sheer scale of adverse events and near-miss incidents that affect certain safety-critical applications. Vincent, Taylor-Adams and Stanhope observe that between 4-17% of patients in acute hospitals suffer from iatrogenic injury [54]. Observational studies have found that 45% of patients experienced some medical mismanagement and 17% suffered events

that led to a longer hospital stay [1]. It has been estimated that approximately 850,000 adverse events occur within the NHS each year [43]. ‘Signals Passed At Danger’ or SPADS provide a further example. These occur when a train passes a signal showing ‘danger’ without authorisation. There were over fifty of these incidents on UK railways in July 2001 [21]. This was nineteen more than in July 2000 but three less than the average for this month over the last six years. Eleven trains ran past the signal by more than 200 yards. One SPAD led to a derailment and two to track damage. In eleven of these cases, the driver had passed a signal at danger before. Since Ladbroke Grove, formal inspections have been required for all SPADS. This has resulted in over 200 formal investigations by Her Majesty’s Railways Inspectorate (HMRI). The nature of these incidents is such that there may also be a far larger number of near-misses. Drivers often manage to avoid passing the signal by rectifying a potential problem ‘at the last minute’ [10]. The proponents of voluntary incident reporting have argued that these near miss incidents provide valuable learning opportunities. They can yield significant insights because they tell us about potential protection mechanisms that might be used in similar circumstances. They also tell us about the potential vulnerability of the system to future incidents in which other drivers might not be so fortunate.

Many organisations have high hopes when they introduce voluntary incident reporting systems. These schemes are not only intended to improve safety by identifying the potential for future failure. It is also, often, hoped that they will reduce costs by avoiding the negative consequences of previous failures. This dual role is illustrated by the objectives that have recently been established for incident reporting within the NHS:

“...the Department of Health should establish groups to work urgently to achieve four specific aims: by 2001, reduce to zero the number of patients dying or being paralysed by maladministered spinal injections (at least 13 such cases have occurred in the last 15 years); by 2005, reduce by 25% the number of instances of negligent harm in the field of obstetrics and gynaecology which result in litigation (currently these account for over 50% of the annual NHS litigation bill); by 2005, reduce by 40% the number of serious errors in the use of prescribed drugs (currently these account for 20% of all clinical negligence litigation); by 2005, reduce to zero the number of suicides by mental health inpatients as a result of hanging from non-collapsible bed or shower curtain rails on wards (currently hanging from these structures is the commonest method of suicide on mental health inpatient wards).” [43]

These are high expectations. Unfortunately, many existing reporting systems have not delivered safety improvements on the scale that some people have predicted. For example, previous paragraphs have described how John Prescott has promoted the expansion of the Scottish CIRAS reporting system to cover the entire rail network. In spite of the perceived success of the CIRAS system, it is hard to demonstrate that Scottish railways have a significantly better safety record than other areas of the network. In June 2001, ScotRail was one of ten train

companies warned by the Railways Inspectorate that it had not done enough to combat the problem of Signals Passed At Danger.

High expectations must also be contrasted with the very prosaic problems that have limited the effectiveness of previous incident reporting systems. For example, it can be difficult to persuade people to contribute reports about near misses or adverse occurrences. The Royal College of Anaesthetist's recent pilot study concluded that the self-reporting of incidents retrieves only about 30% of incidents that can be detected by independent audit. Jha, Kuperman, Teich, Leape, Shea, Rittenberg, Burdick, Segerand, Vander Vliet and Bates confirm this result in their studies of base-line incident frequencies [26]. Their work has detected adverse drug events using three different techniques. Firstly, they have used voluntary incident reporting. Secondly, they have use a computer-based analysis of patient records. Finally, they have performed exhaustive manual comparisons of the same data. In one study, they focused on patients admitted to nine medical and surgical units in an eight-month period [26]. Both the automated system and the chart review strategies were independent and blind. The computer monitoring strategy identified 2,620 incidents. Only 275 were determined to be adverse drug events. The manual review found 398 adverse drug events. Voluntary reporting only detected 23.

3 Fear of Retribution

The Cullen report into the Ladbroke Grove rail accident supported the use of confidential incident reporting as a means of eliciting information about adverse occurrences and near miss events. It was argued that: "(such a system) undoubtedly enables near miss incidents to be reported and receive attention". However, it was also argued that "in the longer term the culture of the industry would be such as to make confidential reporting unnecessary" [10]. Cullen notes that the UK rail industry is some distance away from an ideal situation in which confidential reporting systems would be unnecessary. In most industries, however, employees have a fear of retribution from participation in incident reporting systems. In large-scale national and international systems, this has led to the development of elaborate legal safeguards to protect potential contributors [2]. Unfortunately, most reporting systems lack the funding and the necessary managerial support to provide this level of assurance.

The relatively low participation rates in many reporting systems can be explained by a fear of retribution. In extreme cases, as the Federal Railroad Administration (FRA) notes, employees may even neglect medical treatment rather than expose themselves to workplace harassment:

"FRA has become increasingly aware that many railroad employees fail to disclose their injuries to the railroad or fail to accept reportable treatment from a physician because they wish to avoid potential harassment from management or possible discipline that is sometimes associated with the reporting of such injuries. FRA is also aware that in some

instances supervisory personnel and mid-level managers are urged to engage in practices which may undermine or circumvent the reporting of injuries and illnesses.” [12]

In the medical domain, a number of high-profile cases have acted as a powerful disincentive to the participation in incident reporting systems. For example, the High Court recently intervened to recommend the reinstatement of a surgeon who had expressing worries about the success rate of a colleague in his hospital. The trust initially refused to comply with the Court of Appeal’s finding [6]. This parallels the case of Stephen Bolsin who first uncovered an unusually high death rate among babies undergoing cardiac surgery at the Bristol Royal Infirmary [32]. He subsequently claimed that he was unable to continue working in the NHS as a result of his ‘whistleblowing’ and was forced to move to a hospital at Geelong, near Melbourne. These causes have resulted in the Public Interest Disclosure Act (1998), which allows whistleblowing staff who feel they have been victimised to take their employers to an industrial tribunal. There is no limit to the compensation that can be awarded and employees simply need an “honest and reasonable” suspicion that malpractice has occurred or is likely to occur. Such protection has, however, proven to be insufficient to persuade employees to contribute to many voluntary reporting systems. For example, the 2001 Royal College of Nursing congress explicitly backed a call for action to protect nurses who ‘speak out’. One of the delegates argued that whistleblowing was often seen as ‘grassing up’ or betraying colleagues. Theoretically, such additional protection should not be necessary under the 1998 Act. Some of these concerns can be explained by the informal pressures to conform to the norms of a particular working group. They can also be explained by the practical problems of preserving anonymity within small teams. Given the limited numbers of staff who perform particular tasks on particular shifts, potential contributors can often be identified through a simple process of elimination.

4 Reporting Biases

The previous section has argued that a fear of persecution can prevent staff from participating in incident reporting systems. The success of such schemes in eliciting contributions can also depend upon a number of more complex factors. For example, nursing staff contributed about 90% of all of the reports that have been submitted in a local intensive care unit over the last decade. 621 reports were submitted by nurses compared with 77 reports by medical staff [7]. This suggests that the reporting system may tell us a great deal about the execution of medical procedures. It may, however, tell us relatively little about more complex problems in the planning, coordination and administration of treatment within a department. At one level, it can be argued that this imbalance is due to the reluctance of senior staff to participate in an incident reporting system. However, these figures must be interpreted with great caution. For instance, it is important to consider the total number of staff who might contribute to such a system. Usually the team consisted of three medical staff, one consultant, and up to

eight nurses per shift. The larger number of reports contributed by nursing staff can also be explained in terms of the involvement in, or exposure to, the types of workplace incidents that were solicited under this particular scheme. Nursing staff had the most direct contact with the patients who remain the focus of the reporting system. Hence, it can be argued that they have a proportionately greater opportunity to witness adverse events [7].

Automated logging and tracking systems provide means of addressing the problems both of low contribution rates and of biased participation in a reporting system. The proponents of such systems often have an unfortunate way of advocating their introduction; “competent personnel love them, while incompetent personnel loathe them” [11]. Such assessments hide the difficulty and expense that is often involved in interpreting the data provided by such systems. There is also a concern that any data will be used to punish rather than support staff performance through additional training. These concerns have acted as powerful barriers against the introduction of monitoring equipment onto UK trains. Recommendation 9 of the HMRI report into the accident at Watford South Junction advocated the use of these systems to monitor driving technique. In 1999, however, less than 20% of trains carried this equipment [21]. More recently, the action plan to implement the recommendations of the Southall accident report included steps to extend both voluntary incident reporting systems and automated monitoring equipment [20]. This link between voluntary reporting systems and automated monitoring is instructive. The success of reporting systems in aviation is often explained in terms of the pilot’s fear that any incident may have been observed and reported by their colleagues or by tracking equipment. Participation in reporting system often provides a limited degree of protection and support in any subsequent investigation. However, it is less clear whether such a joint approach might also be extended to healthcare applications. Some initial steps have been taken to use computer-based tools to automatically identify adverse occurrences [26]. Such techniques are inevitably complicated by the difficulty of judging the severity of the patient’s condition prior to treatment.

5 Poor Investigatory Procedures

Even if a reporting system is successful in attracting a large number of submissions, further problems affect the way in which an incident or near-miss is investigated. Theoretical issues complicate the task of determining what should be considered in any investigation. For example, Mackie argues that any event will typically have a number of effects [38]. Any individual is likely only to observe a subset of those effects. This ‘causal field’ is determined by the individual’s ability to observe those effects but also by their prior expectation of what those effects might be. An individual’s interpretation of cause depends upon the subjective frame of reference determined by their causal field. Mackie’s ideas have important consequences because they imply that an investigator’s work may be influenced both by their observations of the effect of an incident and also by their expectations about what those effects will be. For example, if an investigator de-

velops an initial view about the causes of an incident then they may restrict their view of the causal field only to those system behaviours that provide evidence about those causes.

Mackie also argues that many effects stem not simply from a single cause but from a combination of factors that are termed ‘causal complexes’. These ideas are reflected in the UK Health and Safety Executive’s guidance on the incident and accident analysis that support railway safety cases:

“There is much evidence that major accidents are seldom caused by the single direct action (or failure to act) by an individual. There may be many contributing factors that may not be geographically or managerially close to the accident or incident. There might also be environmental factors arising from or giving rise to physical or work-induced pressures. There is often evidence during an investigation that some of the contributory factors have been observed before in events that have been less serious. Accident and incident investigation procedures need to be sufficiently thorough and comprehensive to ensure that the deep-rooted underlying causes are clearly identified and that actions to rectify problems are carried through effectively.” [22]

This quotation also hints at another factor that complicates incident investigations. Statisticians and philosophers, such as Hausman [19], have referred to ‘causal asymmetry’. This embodies the idea that if we know the cause we can predict the likely consequences. However, if we only know the consequences then it is far harder to unambiguously identify a single cause. Typically, many different combinations of events might result in similar consequences. It is precisely this asymmetry that complicates the task of incident investigation and makes it imperative that individual intuition is supported by appropriate investigatory techniques.

A number of more prosaic problems affect the investigation of near-misses and adverse occurrences. There are often insufficient resources to perform a detailed study of the context in which an incident occurred. In local systems, this problem is mitigated by the participant and the safety managers working knowledge of the systems that are described in any contribution. It can, however, be difficult for these local systems to derive independent or expert advice, for example about human factors issues. There is, therefore, a tendency to blame incidents on inadequate attention or on poor staff performance rather looking at the underlying causes of human ‘error’. For example, the reporting systems in a local hospital used acronyms to remind staff to perform particular actions. TAP stood for Tap Aligned Properly. Such advice provides short-term protection against certain classes of adverse events. However, their effectiveness declines rapidly over time. It can also be difficult to ensure that new staff are taught the various incantations that have been proposed. Subsequent study of many of the incidents that helped to generate these acronyms revealed that they were often ‘work arounds’ that were intended to support the use of poorly designed or faulty equipment.

Larger-scale reporting systems can avoid some of these problems by ensuring that their staff are trained in appropriate analytical techniques. Unfortunately, there is little agreement about which approaches might support the causal analysis of incidents in either the rail or the healthcare industries [34, 10]. This lack of consensus has important consequences. It can undermine confidence in the findings of any investigation, especially when there are misgivings about the intent or purpose of any enquiry. HMRI argue that “using the investigation of a SPAD as a means of determining whether, and if so what, disciplinary action should be taken, or as a means of determining questions of liability, for example as between companies, tends to discourage full root cause analysis” [21]

6 Flaws in the Systemic View of Failure

Criticisms of existing practices has prompted the re-training of investigators across the UK rail industry. This has advocated a more ‘systemic’ approach to incident investigation and looks beyond catalytic failures, often characterised by individual human errors, to examine more distal causes. These tend to stem from managerial or organisation failures [45]. This view has been embodied in HMRI guidance that rejects the identification of errors as root causes of incidents on UK railways:

“In these criteria the term ‘root causes’ includes consideration of management’s real and perceived messages to workers, environmental and human factors, as well as plant failures and inadequate procedures. Human errors arising from poor operating conditions, procedures, management expectations or plant design are not root causes; the predisposing factors are.” [22]

Unfortunately, a number of criticisms can be raised against the way in which this systemic view has been interpreted by many safety-critical organisations. Previous sections have argued that adverse occurrences and near miss incidents stem from causal complexes that are difficult to predict. The problems of causal asymmetry also make it difficult to be certain about the precise causes of any mishap. In consequence, many have argued that failures are ‘emergent properties’ that characterise complex, safety-critical systems [36]. Perrow, in particular, has argued that we may have to regard ‘normal accidents’ as the price for technological innovation [44]. A number of objections can be made against this view. The conditions for failure do not suddenly emerge from the application of new technology. Many individuals within safety-critical organisations are often well aware of the potential for an adverse occurrence [46, 53]. In contrast, management pressure to attain other commercial or organisational objectives act to stifle their concerns. Incident reporting can have a limited effect in reiterating the importance of safety concerns, that are typically already known to staff within the organisation. It does not, however, provide a panacea for the deeper organisational issues that have been identified by systemic views of failure.

There is also a deep irony in the systemic view of failure. It starts with the premise that incidents stem from managerial and organisational causes. These create the context in which individual human errors and systems failures can occur. Many managers have, however, argued that the difficulty of predicting these error-inducing contexts helps to absolve organisations from responsibility for particular failures. The implications of the systemic approach to failure can be seen in the words of Daniel Goldin; the head of NASA, when he spoke to the engineers and managers who had been involved in a series of unsuccessful missions to Mars: “As the head of NASA, I accept the responsibility, if anything, the system failed them” [40]. The tension between individual responsibility and the systemic causes of incidents is apparent in this citation. It is also apparent in the behaviour of senior managers within both the rail industry and the health service. At one level, they help to create and control the context in which adverse incidents occur. At another level, they cannot be expected to possess a detailed knowledge of the many different working practices that their staff adopt and that contribute to adverse occurrences. These observations help to explain the current popularity of incident reporting systems; they are seen as a means of communicating safety concerns to higher levels within the management of many organisations. It is, however, far harder to ensure that reporting systems provide accurate information about the potential risks that threaten safe and successful operation. Nor is there any guarantee that higher-levels of management will act on the information that they receive.

7 Analytical Bias

It is important not to underestimate the potential biases that influence the analysis of near misses and adverse occurrences. Over the past three years, we have conducted a series of interviews, surveys and observational studies of incident investigators and safety managers [49]. This work has helped to identify a range of influences that can affect the decision making processes that are intended to distinguish causal factors from the mass of other contextual information that is extracted from an initial report. The following list describes some of these biases. It is not intended to provide an exhaustive account:

Author bias. This arises when individuals are reluctant to accept the findings of any causal analysis that they have not themselves been involved in. For instance, a recent review by the FRA identified that incidents at US highway-rail crossings can trigger investigations by federal organisations, such as the NTSB and the FRA. They can also result in state level enquiries. In some states, responsibility is divided between public agencies and the railroad operators. Elsewhere, responsibility is assigned to regulatory agencies such as the Public Utility Commission, Public Service Commission, or State Corporation Commission. In other states, investigations involve representatives of state, county, and city jurisdictions. Both state and local law enforcement agencies will also be involved if an incident involves the enforcement of traffic laws. Local government bodies are given responsibility for operational matters related to crossings through their or-

dinances. The situation is slightly simpler for incident investigations in the UK. However, railway privatisation has created a situation in which conflict can arise between operating companies, Railtrack and the HMRI. This is neatly encapsulated in Anthony Scrivener’s recent article on Ladbroke Grove entitled ‘Pass the signal - pass the blame’ [47].

Confirmation and Frequency Bias. Confirmation bias arises when investigators attempt to ensure that any causal analysis supports hypotheses that exist before an incident occurs. In other words, the analysis is simply conducted to confirm their initial ideas. Frequency bias occurs when investigators become familiar with particular causal factors because they are observed most often. Any subsequent incident is, therefore, likely to be classified according to one of these common categories irrespective of whether an incident is actually caused by those factors [25]. There are many examples of these two forms of bias in the handling of SPAD reports prior to the Ladbroke Grove accident. Cullen estimates that approximately 85% of all such incidents were classified as the result of driver ‘error’ [10]. The frequency of such findings helped to reinforce this analysis as an acceptable outcome for any SPAD investigation; “I am led to conclude that the ready acceptance of blame by drivers, encouraged by the no blame culture, may have contributed to this poor analysis of root causes”. The subsequent report argued that operating companies should review their incident investigation practices to ensure that there is no presumption that driver error is the sole or principal cause of SPADs.

Recognition bias. This form of bias arises when investigators have a limited vocabulary of causal factors. They actively attempt to make any incident ‘fit’ with one of those factors irrespective of the complexity of the circumstances that characterise the incident. These pressures can be illustrated by the response to initial reports of problems in the performance of cardiac surgery at Bristol Infirmary. The Society of Cardiothoracic Surgeons of Great Britain and Ireland discussed the reports of poor outcomes in 1989. Further information emerged during site visits in 1990. The sub-optimal results were attributed to the low volume of work because an increasing numbers of cases was widely believed to be associated with better outcomes. The eventual enquiry argued that “the focus on throughput may with hindsight be thought to have distracted attention from further inquiry, as the Bristol results, with the exception of the figures for 1990, showed no real improvement” [32].

Political, Sponsor and Professional bias. Political bias arises when a judgement or hypothesis from a high status member commands influence because others respect that status rather than the value of the judgement itself. This can be paraphrased as ‘pressure from above’. Sponsor bias occurs when a causal analysis indirectly affects the prosperity or reputation of the organisation that an investigator manages or is responsible for. This can be paraphrased as ‘pressure from below’. Professional bias arises when an investigators’ colleagues favour particular outcomes from a causal analysis. The investigator may find themselves excluded from professional society if the causal analysis does not sustain particular professional practices. This can be paraphrased as ‘pressure from beside’.

The influence of these workplace issues can be difficult to assess. For example, the FRA Safety Board conducted an analysis of incidents from January 1990 to February 1999. This found that only 18 coded ‘operator fell asleep’ as a causal or contributing factor. The NTSB found these figures difficult to believe given the prevalence of such incidents in other modes of transportation [42]. Two NTSB investigations that had found fatigue as a causal factor were not coded in the FRA database as fatigue-related but as a failure to comply with signals. A number of influences might explain such different interpretations of the same incidents. For instance, the FRA plays a significant role in the promotion of the rail industry as well as in its regulation. The NTSB focuses more narrowly on the investigation of safety-related incidents. In consequences, the political, sponsor and professional influences that act on those organisations will be quite different.

8 Rhetorical Bias and the Problems of Counter-Factual Reasoning

Counter-factual reasoning lies at the heart of most incident investigations [13, 35]. This takes the general form that ‘if a causal factor had not occurred then the incident also would not have taken place’ [37]. If an incident would still have taken place whether or not a event had occurred then it cannot be thought of as causal factor. It is important to stress that counter-factual reasoning is not something new, unusual or surprising. It is often used informally by investigators without realising that this is what they are doing. For example, the US National Transportation Safety Board (NTSB) investigation used this form of reasoning to identify the causal factors in a recent crossing incident; “...had the FRA grade-crossing closure program been more successful in eliminating grade crossings, fewer grade-crossing accidents might have occurred” [41]. Counter-factual reasoning is also used at a more detailed level in the same report; “the train 102 engineer might have seen the long combination vehicle sooner and been able to stop the train in time to avoid the collision if the semitrailer involved had been equipped with retro-reflective tape”. This argumentation style can also help to exclude ‘causes’ that did not contribute to the incident; “..the structural elements of the Northern Indiana Commuter Transportation District railcar 11 collision post that failed were overwhelmed by the force of the collision, and the post could not have prevented penetration of the steel coil, given the train speed and the weight of the coil” [41].

It is difficult to underestimate the prevalence of counter-factual reasoning in the analysis of adverse occurrences and near-miss incidents. It forms a major component of the techniques advocated by NASA [39] and the US Safety Systems Society [50]. Counter-factual reasoning can, however, pose numerous problems for incident investigation. For example, how sure can we be that an incident would not have occurred if a causal factor had not been present? Causal asymmetries suggest that many different causal complexes will have the same outcome. For instance, there are few guarantees in the previous incident that the engineer would have been able to avoid the collision even if the semi-trailor had

been equipped with reflective tape. The problems of inattention and fatigue in previous incidents have shown that such safeguards do not provide guaranteed protection against adverse occurrences. The complex issues surrounding counterfactual reasoning is a research area in its own right. Byrne and her colleagues have conducted a number of preliminary studies that investigate the particular effects that characterise individual reasoning with counterfactuals [8, 9]. This work argues that deductions from counterfactual conditionals differ systematically from factual conditionals and that, by extension, deductions from counterfactual disjunctions differ systematically from factual disjunctions. This is best explained by an example. If we argue that “the train 102 engineer might have seen the long combination vehicle sooner and been able to stop the train in time to avoid the collision if the semitrailer involved had been equipped with retro-reflective tape” readers will infer that the semitrailer was not, in fact, equipped with retro-reflective tape. This counter-factual style of argument can have such a persuasive effect that readers overlook contradictory evidence elsewhere in a report [31]. There are more complex examples of the inferences that readers draw from counter-factual arguments. The statement that ‘either the brakes were applied too late or the train was going too fast’ is a factual disjunction. Byrne argues that such sentences encourage the reader to think about these possible events and decide which is the most likely. There is an implication that at least one of them took place. The statement that ‘had the FRA grade-crossing closure program been more successful or the semitrailer been equipped with retro-reflective tape then the incident would have been avoided’ is a counterfactual disjunction. Byrne argues that this use of the subjunctive mood not only communicates information about the possible outcome of the incident but also a presumption that neither of these events actually occurred.

This theoretical work has pragmatic implications for incident investigation. If factual disjunctions are used then care must be taken to ensure that one of the disjuncts has occurred. If counter-factual disjunctions are used then readers may assume that neither disjunct has occurred. The distinction between counterfactual and factual disjunctions forms part of a wider concern to ensure that analytical biases are not hidden through the inappropriate use of language in incident reports. For example, rhetorical devices known as tropes can be used to increase the impact and effectiveness of everyday prose. They can also be used to achieve particular effects on the readers of an incident report. The following paragraphs provide a brief introduction to the techniques that have been used within the rail and healthcare industries. A more sustained analysis is presented in [29].

Amplification involves the restatement of an idea or argument. It often also involves the introduction of additional details. For example, the US Food and Drugs Administration (FDA) recently described actions that were taken in response to incidents involving nutritional supplements:

“The recognition of a cluster of cases was the key to the detecting of Eosinophilia-Myalgia Syndrome (EMS). Interactions among various spe-

cialists, including a family physician, hematologist, rheumatologist, clinical immunologist and epidemiologists, was crucial to this process.

Both EMS's clinical seriousness, and uncertainties surrounding its etiology, indicate the need for health professionals to remain vigilant regarding adverse events possibly associated with the use of L-tryptophan-containing dietary supplements, and to report such events to MedWatch." [14]

The first paragraph notes the importance of detecting the initial cluster of incidents using a reporting system. The second paragraph amplifies this by noting the importance of the MedWatch reporting system as a means of detecting adverse events that might be associated with L-tryptophan in dietary supplements. This technique can have the effect of drawing the reader's attention to a particular concept or idea. The amplification not only introduces new facts but it also supports and reiterates the arguments that are introduced in previous sentences. This technique can create problems when the amplification of particular aspects of a previous assertion can detract from other arguments or items of information. It is, therefore, important to establish the credibility of both an initial assertion and the subsequent amplification. For instance, an article about the cluster in the Albuquerque Journal News on the 7th November helped to trigger the FDA release of the public advisory on the 11th November. The initial cluster was not triggered by submissions to the reporting system, as might have been inferred from the use of amplification in the previous example.

Anaphora uses repetition at the beginning of successive phrases, clauses or sentences. It can create an impression of climax in which the repetition leads to a particularly important insight or conclusion.

"In April 1990, two more cases of sudden death associated with the use of barium enema kits were reported. A 41 year-old female complained of nausea shortly after insertion and inflation of the tip/cuff assembly, went into cardiac arrest within 30 seconds and underwent unsuccessful resuscitation efforts. In the third case, a 72 year-old female had an immediate reaction after the tip portion of the tip/cuff assembly was inserted prior to introduction of the barium contrast agent, went into vascular collapse and died." [14]

This example illustrates the successive use of the phrase 'A XX-year-old female' to build up descriptions of similar incidents. The investigator uses each successive sentence to 'stack up' evidence in a manner that will eventually support their analysis of common causes. It is important to emphasise that such techniques are not of themselves either 'good' or 'bad'. Rhetorical devices can be used to convince us of well-justified conclusions or to support half-baked theories. It is important, however, to be sensitive to the effects that such techniques might have on the readers of an incident report. For instance, the previous citation can be interpreted to provide readers with a clear summary of the evidence that supports the investigators' conclusions. It can also be interpreted in a more negative

light. The repetition of such phrases may create an impression of certainty about the similarities between incidents that might not be justified by the evidence.

Antithesis uses juxtaposition to contrast two ideas or concepts. This can be illustrated by the use of the terms ‘properly’ and ‘improperly’ in the following report into a needle-stick incident:

“The housekeeper reported that he had used extra force to push the lid down, because the container was overfilled... Another housekeeper stuck her finger while removing a full sharps container from a wall bracket. . . When used properly, sharps containers can prevent needle-stick injuries. When used improperly, they can create a serious hazard.” [15]

Here the consequences of ‘proper use’ are contrasted with those of ‘improper use’. This technique is important because readers may make a number of additional inferences based upon such constructions. In this context, it is tempting to infer that the preceding injuries were sustained as the result of improper use, although this is not explicitly stated. It is also important to consider that the author presents no evidence to support their assertions about the consequences of proper and improper use. This is important because the citation emphasizes that idea that violation of proper procedures will result in injury. It does not consider that poor design might have created the potential for such serious consequences.

Asyndeton omits conjunctions between words, phrases and clauses. This technique creates an impression of ‘unpremeditated multiplicity’ [18]. The reporter can think of so many elements in the list that they hardly have time to introduce explicit conjunctions. There is also a sense in which this technique builds to a particular conclusion. This is illustrated by the opening sentences of a report into medical error; “a 62-year-old man came to the Emergency Department complaining of mid-sternal pain, shortness of breath, diaphoresis, and nausea after shoveling snow” [16]. *Asyndeton* creates precise and concise summaries. It, therefore, offers considerable stylistic benefits to more verbose explanations. There are, however, dangers. For instance, the use of such enumerations can create an impression of completeness where none was intended. There may have been other indications about the patient’s condition that were omitted from the list. There is, however, a strong tendency for readers to regard the enumeration as complete unless the final conjunct is omitted; ‘a 62-year-old man came to the Emergency Department complaining of mid-sternal pain, shortness of breath, diaphoresis, nausea after shoveling snow’. This creates a converse problem for the reader. The implied omission of closing conjuncts, as in the previous example, can lead to uncertainty about the information that might have been omitted from the list.

Conduplicato relies upon the repetition of key words or phrases at, or very near the beginning, of subsequent sentences. *Conduplicato* provides a focusing device because writers can use it to emphasise key features in preceding sentences. This helps to ensure that readers notice concepts or ideas that may have been overlooked when they read the initial sentence. This can be illustrated by the following quotation in which contributors stress the importance of “turning off” the ventilation sensors

“Medical intervention was needed to turn off the minute ventilation sensor in each pacemaker. When the sensors were turned off, the patients’ heart rates returned to normal.” [16]

Most contributors and safety managers draft incident reports without ever being aware that they are exploiting such rhetorical devices. They inadvertently construct prose that supports their arguments without explicitly considering the impact that their use of language will have upon their readers. They may inadvertently stress conclusions that are not well supported by the available evidence. They may also cast doubt on other findings that contradict their version of events. Unfortunately, this inadvertent use of rhetorical devices is often exposed at litigation. This should not be surprising. Many law courses, especially in the United States, include training in the use of rhetoric.

9 Classification Problems

Many regulatory and investigatory organisations have begun to codify information about previous incidents. This has numerous benefits. Firstly, the use of such codes can help to strip out the rhetorical effects that bias the interpretation of natural language accounts. Secondly, classification schemes provide key terms that can be used to access incident data in large, national and international databases. Unfortunately, a number of problems affect the practical application of this approach. If the codification of incidents is performed centrally then it is important that staff understand enough about the context in which an incident occurs for them to ensure that the correct codes are assigned. Alternatively, if incidents are to be codified at a local level then it can be difficult to ensure that different safety managers assign the same codes to similar incidents. For example, the FDA describe a case study in which a violent patient in a wheelchair was suffocated through the use of a vest restraint that was too small. The risk manager, JC, proceeded as follows:

“She finds the list of event terms, which was detached from the rest of the coding manual... She muses: ‘Mr. Dunbar had OBS which isn’t listed in these codes; he had an amputation which is listed; he had diabetes which isn’t listed; and he had hypertension which is listed’. JC promptly enters 1702 (amputation) and 1908 (hypertension) in the patient codes. She then finds the list for Device-Related Terms... She reviews the terms, decides there was nothing wrong with the wheelchair or the vest restraint, and leaves the device code area blank.” [17]

The resulting classification of 1702 (amputation) and 1908 (hypertension) provided few insights into the nature of the incident. This classification is more misleading than corresponding prose accounts even considering the potential biases that can be introduced through rhetorical effects and poorly constructed counter-factual arguments. The problems of incident classification also affect the rail industry. SPAD investigation procedures require that each incident is assigned a causal category. A recent HMRI report described how:

“Inspectors found there were difficulties with two of these categories: ‘misjudgement’ and ‘disregard’. Although miscategorisation is not thought to be widespread, examples were found of SPAD incidents which had been categorised ‘disregard’, which actually seemed to be instances of driver misjudgement. An example was seen in RTMZ where a Virgin Trains driver had appeared to make every effort to brake at Coventry signal CY37 in poor weather conditions, yet the incident was categorised as ‘disregard’ rather than ‘misjudgement’. Inappropriate categorisation should be avoided, otherwise it will reduce the credibility of the SPAD incident statistics and affect the rating of specially monitored drivers.” [21]

A number of problems remain to be addressed even if incidents can be ‘correctly’ classified according to appropriate taxonomies. The most important of these relates to the storing and retrieval of large collections of codified incident reports. At present, incident reporting systems rely upon relational database technology. Each incident is classified according to a number of pre-determined fields. This approach has a number of consequences. It can lead to an extremely static classification system. There is no automatic means of reclassifying thousands of previous incidents if changes are made to a taxonomy. Some reporting schemes now hold more than 500,000 reports [2]. The scale of such systems creates problems if we must reclassify historical data to reflect changes in the coding scheme. If such changes are not made then there is a danger that safety managers may fail to discern that recent incidents form part of a wider pattern, which is obscured by weaknesses in the previous classification scheme. This problem is particularly acute when taxonomies are extended to describe human behaviour, as in the previous citation. The field of human factors research has changed rapidly over the last decade with an increasing focus on group interaction. However, few of these changes have been reflected in incident reporting systems because of the costs associated with manually analysing and re-classifying existing records.

Further problems affect the use of relational databases. The theoretical underpinnings of these systems are often poorly understood by the people who must use them. Safety managers, therefore, often rely upon pre-formulated queries to sort, filter and combine incident data. These queries are pre-programmed by system administrators who typically have a clearer understanding of the semantics of the commands that are being issued to the incident database. A consequence of this is that safety managers are often not being provided with the information that they think they are requesting each time they issue a query. It can also be difficult for safety managers to formulate the queries that they really would like to ask of their system because they lack the necessary technical knowledge about the implementation and operation of relational database technology [28].

We have experimented with a number of alternative technologies that address the problems, described above. Probabilistic information retrieval and conversational case based reasoning systems enable users to search for data without forming complex, structured queries [28]. A further benefit is that web-based search techniques can help to automate the indexing of large collections of inci-

dent reports. This avoids the overheads associated with the manual classification that may be necessary when changes are made to the underlying models that structure relational databases. Unfortunately, probabilistic information retrieval and conversational case based reasoning tools suffer from other problems. In particular, it can be difficult to ensure that particular queries yield appropriate levels of precision and recall. A system can exhibit poor precision if it returns many incidents that the user does not believe are related to their query. The user must then manually filter the large number of incidents that the system considers to be a match. Conversely, poor recall occurs when a system fails to return an incident that the user believes is related to their query. This can prevent analysts from determining that an incident forms part of a wider pattern. These problems are compounded by the computational relationship between precision and recall. Systems that provide good recall are often imprecise. Conversely, systems that offer high degrees of precision will often exclude incidents that ought to have been returned as a potential match.

10 Reliance on Reminders

The success of any reporting system depends on the interventions that are triggered by the information that it provides. There are many notable success stories where agencies have responded in a prompt and effective manner [15, 16]. Equally, however, a number of limitations constrain the use of incident data to inform safety management. Incident reporting systems often yield few surprises. The organisations that establish and operate these systems already have a good idea of the safety issues that affect their working practices. This can be illustrated by pioneering research into incident reporting within NHS hospitals in the North-West of England [52]. One study focussed on 19 incidents that were reported over approximately one month to an Accident and Emergency Department. The particular analysis technique that was applied to this data yielded a total of 93 potential causes. 45% of these related to organisational issues while 41% were classified as ‘direct’ human causes. The organisational causes included the need to secure external services. In particular, incidents were often triggered or exacerbated by the need to secure beds for the patients in the Department. They also included a lack of senior staff during peak periods. Direct human causes included problems that new Senior House Officers experienced in interpreting X rays. They also stemmed from a culture of learning from mistakes and a reluctance to contact senior staff. A similar study was then conducted into incident reporting within an Anaesthesia Department. This yielded 15 incidents with 78 root causes. 27% of these were identified as organisational issues, 40% stemmed from direct human causes and 26% were related to technical issues. The incidents were argued to illustrate less effective protocols than had been established in the Accident and Emergency Department. Several incidents indicated confusion over which drugs to stock and when. The study also revealed design problems with particular devices and the inadequate training that some staff received before being required to operate new systems. It is ironic

that most of these issues had already been identified as significant problems by safety-managers within the hospitals. Such concerns do not, however, secure the resources and wider organisational support that is necessary to address many of these issues. This problem can be illustrated by a recent conversation with an NHS trust manager. He argued that incident reporting saves money by avoiding litigation and by reducing the amount of time that a patient might otherwise have to spend in hospital. Incident reporting does not, however, generate the resources that are needed to invest in addressing safety-related problems.

The impact of these financial issues on the effectiveness of incident reporting systems cannot be underestimated. The most frequent remedial action in one intensive care unit within an NHS hospital was to disseminate staff reminder statements. In the period from August 1995 to November 1998, 82 ‘Remind Staff?’ statements were issued out of a total of 111 recommendations [7]. The 29 other recommendations concerned the creation of new procedures or changes to existing protocols (e.g. ‘produce guidelines for care of arterial lines - particularly for femoral artery lines post coiling’), or were equipment related (e.g. ‘Obtain spare helium cylinder for aortic pump to be kept in ICU’). None of the recommendations addressed the organisational or managerial issues that have been identified as a potential target for incident reporting systems and which are the focus for systemic views of failure. Such issues were beyond the scope of the system. This is not an isolated example. Similar patterns can be identified within aviation reporting systems [27]. The reiteration of well-known safety recommendations raises fundamental concerns about the utility of incident reporting systems. Human factors research points to the dangers of any reliance on reminders. Unless people are continually reminded then they are likely to forget the importance of safety precautions over time [24].

11 Inadequate Risk Assessments

There are many reasons why the recommendations that are derived from incident reporting systems often rely upon ‘short-term fixes’ rather than addressing the underlying causes of incidents and accidents. As mentioned above, there is a perception that they are a cost-saving measure. Reporting systems are, therefore, isolated from the revenue streams that might otherwise support necessary investments. Incident reporting schemes are also poorly integrated into wider forms of risk assessment. I recently witnessed the bizarre situation in which a design team were using Bayesian techniques to derive best estimates for reliability data while others, in the same organisation, had numerical data for the same faults [30].

Some incident reporting systems do not conduct any formal risk assessment for the near misses and adverse occurrences that they identify. This creates problems for anyone who wants to exploit the incident data. There is no way of distinguishing whether one incident must be addressed before another even though one might have a relatively high probability of recurrence and the potential for severe adverse consequences [29]. In other organisations, groups have

drawn attention to such high-risk incidents without provoking an appropriate response:

“During the almost five years preceding the Ladbroke Grove accident, there had been at least three occasions when some form of risk assessment analysis on the signaling in the Ladbroke Grove area has been suggested or proposed. The requests were: the Head of Technical Division’s letter of 11 November 1996 which requested a layout risk assessment of the re-signaling (paragraph 43); the Field Inspector’s letter of 16 March 1998 to Railtrack (paragraph 64); and the Railtrack Formal Inquiry of 1 July 1998 (paragraph 66). In addition there was an earlier request for details of measures taken to reduce the level of SPADs in the area around SN109 recorded in the Head of Technical Division’s letter of 1st March 1995 (paragraph 39). None of these requests appear to have been pursued effectively by HMRI.” [23]

Even when risk assessments are performed, there can be biases that emerge in the criteria that are used. This is often unavoidable. For example, risk assessment typically involves some appraisal of the frequency and consequence of an event. However, with a near-miss incident one cannot assume that any future recurrence will have the same outcome. Many reporting systems, therefore, assume the ‘plausible worst case scenario’ is an approximation to the potential consequences of a failure. Some of the problems of assessing the consequences of adverse incidents can be identified in W.S. Atkins’ recent report into the investigation of SPADs on UK railways:

“The system for selection of incidents for full investigation is skewed towards shunting incidents, which are often of low consequence (both actual and potential). This arises from the simplicity of the severity category system used to provide an initial classification of the seriousness of incidents. In part the system classifies by length of overrun, an approach that we consider to be unhelpful because an incident at a shunt signal with zero overlap results in a high rating which is often out of proportion to its seriousness. A long overrun on plain line (maybe past the next signal at green) is often accorded undue seriousness. Disregard of a cautionary aspect, followed by a very short SPAD (with potentially serious consequences had the brake application occurred fractionally later) receives a disturbingly modest rating. A further weakness with the severity category system is that no distinction is drawn between derailments on running lines and contained derailments on trap points. This is distinctly unhelpful because in the latter case it is liable to result in inappropriate amounts of attention and effort being focussed on intrinsically low consequence events.” [48]

This quotation illustrates an important barrier to the successful implementation of incident reporting systems. Unless they are supported by clear guidelines to help assess the potential frequency and consequence of any recurrence there

is little likelihood that management will allocate sufficient resources to rectify major safety problems. As we have seen, consequence estimates are error-prone and difficult to validate. Previous sections have also described the problems that arise in deriving accurate estimates of incident frequencies from reporting systems. Many incidents are not reported because of a fear of retribution. Automated logging systems provide a greater assurance that potential failures will be detected, however, a recent NTSB symposium identified numerous instances in which these systems failed to provide reliable data [3].

12 Conclusion and Further Work

The last three years has seen a rapid growth in the number and scale of incident reporting systems. The Ladbroke Grove accident stimulated a range of initiatives in the UK rail industry, including John Prescott's expansion of the CIRAS reporting system [10]. The Bristol Infirmary enquiry has had a similar impact on the UK healthcare industry [32]. Strong claims have been made about the potential benefits of these systems. Incident reporting applications are perceived to offer valuable insights into the near-miss incidents that have the potential to threaten future safety. They can also be used to elicit information about 'lessons learned' and act as an exchange for best practice [52].

This paper has, however, argued that significant barriers must be addressed before incident reporting systems can be successfully applied within many industries. These can be summarised as follows:

1. *unrealistic expectations.*

Many people who initiate reporting systems expect reductions in the frequency and consequence of adverse events that are unreasonable given previous experience in running these schemes. These expectations are particularly problematic given that many types of incident will not be reported to confidential systems. There can be strong organisation and cultural barriers that prevent employees from disclosing information about their friends and colleagues;

2. *fear of retribution.*

Some local systems enjoy good levels of participation while trusted individuals administer the scheme. Staff learn to trust the integrity of those individuals. However, when they are replaced participation rates may fall dramatically [7]. This effect is clearly linked to potential contributors' concerns that they will be viewed as 'whistleblowers' either by their colleagues or by those who administer the system.

3. *reporting biases.*

Even once confidence has been established in a system, there are few guarantees that all staff will contribute incident reports. Variations in participation rates have been observed both within working groups at the same location, as in hospital systems, and between geographical regions, for example across

the rail network. Automated systems are increasingly being introduced to trigger investigations into near-miss incidents. However, some tasks cannot easily be instrumented. Many of the more specialised monitoring systems are unreliable and often provide ‘false positives’ that consume finite analytical resources. In consequence, it seems likely that reporting rates of less than 20-30% will be typical of many healthcare applications. These problems do not affect some reporting systems. SPAD reports provide a relatively accurate impression of the frequency of these events. However, the monitoring systems that help to detect these incidents tell us very little about incidents that *almost* resulted in a SPAD but that were narrowly averted by operator intervention.

4. *poor investigatory procedures.*

Once an adverse occurrence or near miss has been reported, it can be difficult to determine what factors should be included within an investigation. This is important for theoretical reasons because it can be difficult to identify salient factors within what Mackie terms the ‘causal field’ [38]. Hausman also points to the problems created by ‘causal asymmetry’ [19]. If we know the cause then we can determine the effects, however, if all we observe are the effects then it can be difficult to reach firm conclusions about the multiple possible causes of those effects. These theoretical problems are exacerbated by the resource constraints that affect incident reporting. Many organisations lack both the funding and the expertise to investigate more than a single causal hypothesis. This clearly limits the value of any insights that might be obtained from the analysis of near miss incidents.

5. *flawed systemic views of failure.*

The limited resources that are available to fund the analysis of many incidents are stretched by the recent emphasis on the systemic causes of failure. The proponents of this approach have urged investigators to look beyond the catalytic, triggering events of individual human error to look at the deeper systemic causes that are often related to organisational and managerial issues [36]. Many reporting systems avoid the practical and ‘political’ difficulties that such studies entail by limiting the scope of their analysis. Some of the local systems ‘target the doable’ [7]. There have also been examples of organisations whose upper levels of management have actively exploited arguments about the systemic causes of failure to mitigate managerial responsibility for particular failures. They have argued that the ‘system’ failed them in a manner that could not have been predicted before the failure. This stretches the interpretation of the systemic view of failure; it has also created a situation in which organisations apparently accept that it may not be possible to use any form of analysis, including incident reporting, to anticipate and respond to future failures.

6. *analytical bias.*

There are numerous forms of bias that can affect the analysis of incidents

once they have been reported. We have briefly described author bias, judgement and hindsight bias, confirmation and frequency bias, recognition bias, political, sponsor and professional bias. This is not an exhaustive list but it illustrates the difficulty of ensuring that any investigation is not hindered by ‘undue’ influences. These issues are particularly important in incident reporting when many stages of an initial investigation and analysis will be performed not by an external authority but by the organisation that was directly involved in the occurrence.

7. *rhetorical bias and the problems of counter-factual reasoning.*

A variety of rhetorical techniques can be used to ‘hide’ analytical bias within an incident report. It can be difficult to avoid using these tropes when constructing prose arguments to support particular findings. It is, therefore, important that both the readers and the writers of these reports are sensitive to the effects that these techniques can have. For instance, the repeated reference to particular items of evidence can indirectly increase the salience of that information. The psychological effects of rhetorical devices are mirrored by the unintended inferences that can be drawn from counter-factual reasoning. This style of argument takes the following form; ‘X is a causal factor if the incident would not have occurred if X also had not occurred’. As mentioned, counter-factual reasoning can be fraught with dangers. There is often an implicit and unwarranted assumption that X did, indeed, occur [19].

8. *classification problems.*

Many organisations have responded to the problems of interpreting prose descriptions by adopting causal taxonomies. These initiatives form part of a wider attempt to classify incidents according to a range of different criteria. This offers numerous benefits. In particular, the elements of the classification be used as indexing terms in relational databases. Unfortunately, field studies have shown that few safety managers know how to use these tools to accurately extract information about previous incidents. Problems also arise when the items in a database have to be manually reclassified to reflect changes in a causal taxonomy. This can be particularly onerous for national systems that hold many hundreds of thousands of records. Several prototype systems have been developed to address these problems. For instance, we are using information retrieval techniques that were originally developed for mass-market web-based applications. These approaches are the subject of on-going research and currently suffer from poor precision and recall.

9. *reliance on reminders.*

Many reporting systems lack the financial resources that are necessary to address underlying system failures. These systems are, typically, seen as a form of cost reduction rather than as a form of income generation. This separation of reporting systems from sources of investment can result in recommendations that focus narrowly on ‘quick fixes’. Studies of previous systems have seen a tendency to adopt a perfective approach in which operators are urged

to try harder to avoid future incidents. Such reminder statements provide dubious protection given that they must be continually reinforced if they are not to be forgotten.

10. *inadequate risk assessment.*

The design of safety-critical applications is typically guided by some form of risk assessment. Risk can be thought of as the product of the consequence and the likelihood of a particular failure. Incident reporting systems have been proposed as powerful means of informing risk assessments. They can provide quantitative data about the relative frequency of previous failures [34]. As we have seen, however, analytical and reporting biases undermine such statements. Similarly, the nature of ‘near miss’ incidents makes it very difficult to identify the ‘plausible worst case scenario’ that might inform any decision about the consequences of a future recurrence.

This is a partial list. For instance, we have not considered the powerful influence that a fear of media publicity can have upon the dissemination of safety-related information about previous mishaps. Similarly, previous paragraphs have not mentioned the conflicts that can arise when external incident reporting agencies must rely upon funding approval from the managers of the organisation that they collect reports about. A more complete introduction to the problems of incident reporting and a detailed explanation of potential solutions are presented in a forthcoming *Handbook of Incident Reporting* [29].

We have illustrated the problems that frustrate incident reporting using examples drawn from existing systems in the rail and healthcare industries. Many of these applications, especially within the healthcare industry, have been sponsored by individuals with a personal motivation for identifying safety issues in their workplace. Other systems operate on a far larger scale, such as the SPAD reporting process for UK rail operators. As we have seen, however, many of these diverse systems have faced remarkably similar problems. For instance, it is difficult to validate the findings of any causal analysis. It can also be difficult to assess the risks of future recurrences. It is regrettable that the proponents of recent initiatives to set-up national incident reporting systems have not taken more time to consider the range of technical problems that complicate the operation of these existing systems.

References

1. L.B. Andrews, C. Stocking, T. Krizek, L. Gottlieb, C. Krizek, and T. Vargish. An alternative strategy for studying adverse events in medical care. *Lancet*, (349):309–313, 1997.
2. Aviation Safety Reporting System. The Aviation Safety Reporting System. Technical report, NASA Ames Research Centre, California, United States of America, 2000. <http://asrs.arc.nasa.gov>.
3. K. Bolte, L. Jackson, V. Roberts, and S. McComb. Accident reconstruction/simulation with event recorders. In *International Symposium on Transporta-*

- tion Recorders*, pages 367–369. National Transportation Safety Board, Washington DC, USA, 1999. <http://www.nts.gov/publicn/1999/rp9901.pdf>.
4. British Broadcasting Corporation. Rail summit moves forward on safety. Technical report, News Staff, BBC, London, United Kingdom, 30th November 1999. http://news.bbc.co.uk/hi/english/uk/newsid_543000/543019.stm.
 5. British Broadcasting Corporation. Plan to stop dangerous doctors. Technical report, News Staff, BBC, London, United Kingdom, 13 June 2000. http://news.bbc.co.uk/hi/english/health/newsid_788000/78805.stm.
 6. British Broadcasting Corporation. Doctors back down in whistle-blower case. Technical report, News Staff, BBC, London, United Kingdom, 2001. http://news.bbc.co.uk/hi/english/health/newsid_1470000/1470590.stm.
 7. D. K. Busse and D. J. Wright. Classification and analysis of incidents in complex, medical environments. *Topics in Health Information Management*, 20(4):1–11, 2000. Special Edition on Human Error and Clinical Systems.
 8. R.M.J. Byrne and S.J. Handley. Reasoning strategies for suppositional deductions. *Cognition*, pages 1–49, 1997.
 9. R.M.J. Byrne and A. Tasso. Deductive reasoning with factual, possible and counterfactual conditionals. *Memory and Cognition*, pages 726–740, 1999.
 10. Cullen. *The Ladbroke Grove Rail Inquiry Part 1 Report*. HSE/Stationary Office, London, United Kingdom, 2001. <http://www.hse.gov.uk/railway/padtrail/lgril.pdf>.
 11. M. Durkin. Digital audio recorders: Life savers, educators and vindicators. In *International Symposium on Transportation Recorders*, pages 139–144. National Transportation Safety Board, Washington DC, USA, 1999.
 12. Federal Railroad Administration. Fra guide for preparing accidents/incidents reports. Technical Report DOT/FRA/RRS-22 Effective: January 1997, Office of Safety, Federal Railroad Administration, Washington DC, United States of America, 1997. <http://safetydata.fra.dot.gov/Objects/guide97.pdf>.
 13. T. S. Ferry. *Modern Accident Investigation and Analysis*. John Wiley and Sons Inc., London, 1988.
 14. Food and Drug Administration. Clinical impact of adverse event reporting. Technical report, Department of Health and Human Services, Public Health Service, US Food and Drug Administration, Rockville, Maryland, USA, 1996. <http://www.fda.gov/medwatch/articles/medcont/synopses.htm>.
 15. Food and Drug Administration: A. Morrison. Avoiding sticks from sharp containers. *User Facility Reporting Bulletins*, 1998. <http://www.fda.gov/cdrh/fusenews/fuse25.pdf>.
 16. Food and Drug Administration: D. Dwyer. Sending the wrong signals. *User Facility Reporting Bulletins*, 2000. <http://www.fda.gov/cdrh/fusenews/ufb33.html>.
 17. Food and Drug Administration: M. Weick-Brady. Those codes! *User Facility Reporting Bulletins*, 1996. <http://www.fda.gov/cdrh/issue18.pdf>.
 18. R. Harris. A handbook of rhetorical devices. Technical report, SCC, Cosa Mesa, California, 1997. http://www.sccu.edu/faculty/R_Harris/rhetoric.htm.
 19. D.M. Hausman. *Causal Asymmetries*. Cambridge University Press, Cambridge, U.K., 1998.
 20. Health and Safety Commission. The Southall Rail Accident Inquiry Report: HSC action plan to implement recommendations. Technical report, Health and Safety Executive, London, United Kingdom, 2000. <http://www.hse.gov.uk/hsc/south-01.htm>.

21. Her Majesty's Railway Inspectorate. Report on the inspection carried out by hm railway inspectorate during 1998/99 of the management systems in the railway industry covering signals passed at danger. Technical report, Health and Safety Executive, London, United Kingdom, 1999. <http://www.hse.gov.uk/railway/spad-01.htm>.
22. Her Majesty's Railway Inspectorate. Assessment criteria for railway safety cases. Technical report, Health and Safety Executive, London, United Kingdom, 2000. <http://www.hse.gov.uk/railway/criteria/index.htm>.
23. Her Majesty's Railway Inspectorate. Internal inquiry report: Events leading up to the ladbroke grove rail accident on 5 october 1999. Technical report, Health and Safety Executive, London, United Kingdom, 2000. <http://www.hse.gov.uk/railway/paddrail/inq-03.htm>.
24. V.D. Hopkin. The impact of automation. In M.W. Smolensky and E.S. Stein, editors, *Human Factors in Air Traffic Control*, pages 391–419. Academic Press, London, United Kingdom, 1998.
25. D. Javaux. The cognitive complexity of pilot-mode interaction. In *HCI-Aero'98: Conference on Human-machine Interaction in Aeronautics*, Montreal, Canada, 1998.
26. A.K. Jha, G.J. Kuperman, J.M. Teich, L. Leape, B. Shea, E. Rittenberg, E. Burdick, D.L. Seger, M. Vander Vliet, and D.W. Bates. Identifying adverse drug events: development of a computer-based monitor and comparison with chart review and stimulated voluntary report. *Journal of the American Medical Informatics Association*, 5(3):305–314, 1998.
27. C.W. Johnson. Don't keep reminding me: The limitations of incident reporting. In K. Abbott, J.-J. Speyer, and G.Boy, editors, *HCI Aero 2000: International Conference on Human-Computer Interfaces in Aeronautics*, pages 17–22, Toulouse, France, 2000. Cepadues-Editions.
28. C.W. Johnson. Software support for incident reporting systems in safety-critical applications. In F. Koornneef and M. van der Meulen, editors, *Computer Safety, Reliability and Security: Proceedings of 19th International Conference SAFECOMP 2000*, LNCS 1943, pages 96–106. Springer Verlag, 2000.
29. C.W. Johnson. *A Handbook of Incident Reporting: A Guide to the Detection, Mitigation and Avoidance of Failure in Safety-Critical Systems*. Springer Verlag, UK, 2001 (in press).
30. C.W. Johnson, G. Le Galo, and M. Blaize. Guidelines for the development of occurrence reporting systems in european air traffic control. Technical report, European Organisation for Air Traffic Control (EUROCONTROL), Brussels, Belgium, 2000.
31. C.W. Johnson, J.C. McCarthy, and P.C. Wright. Using a formal language to support natural language in accident reports. *Ergonomics*, 38(6):1265–1283, 1995.
32. I. Kennedy. *Learning from Bristol: the report of the public inquiry into children's heart surgery at the Bristol Royal Infirmary 1984 -1995*. Command Paper: CM 5207. Her Majesty's Stationary Office, London, United Kingdom, 2001. <http://www.bristol-inquiry.org.uk>.
33. U. Kjellen. *Prevention of Accidents Through Experience Feedback*. Taylor and Francis, London, United Kingdom, 2000.
34. L. Kohn, J. Corrigan, and M. Donaldson. *To Err Is Human: Building a Safer Health System*. Institute of Medicine, National Academy Press, Washington DC, United States of America, 1999. Committee on Quality of Health Care in America.
35. P.B. Ladkin. Causal reasoning about accidents. In F. Koornneef and M. van der Meulen, editors, *SAFECOMP 2000*, Lecture Notes in Computing Science No. 1943, pages 344–355. Springer Verlag, Berlin, Germany, 2000.

36. N.G. Leveson. *Safeware: System Safety and Computers*. Addison Wesley, Reading, MA, United States of America, 1995.
37. D. Lewis. *Counterfactuals*. Oxford University Press, Oxford, UK, 1973.
38. J.L. Mackie. Causation and conditions. In E. Sosa, editor, *Causation and Conditions*. Oxford University Press, Oxford, 1975.
39. NASA. NASA procedures and guidelines for mishap reporting, investigating and record-keeping. Technical Report NASA PG 8621.1, Safety and Risk Management Division, NASA Headquarters, Washington DC, USA, 2001. <http://www.hq.nasa.gov/office/codeq/doctree/safeheal.htm>.
40. NASA (D. Goldin). "When The Best Must Do Even Better" Remarks by NASA Administrator Daniel S. Goldin At the Jet Propulsion Laboratory Pasadena, CA March 29, 2000. Technical report, NASA Headquarters, Washington DC, USA, 2000. http://www.hq.nasa.gov/office/pao/ftp/Goldin/00text/jpl_remarks.txt.
41. National Transportation Safety Board. Railroad Accident Report Collision of Northern Indiana Commuter Transportation District Train 102 with a Tractor-Trailer Portage, Indiana June 18, 1998. Technical Report NTSB/RAR-99/03, NTSB, Washington, DC United States of America, 1999. <http://www.nts.gov/Publictn/1999/RAR9903.pdf>.
42. National Transportation Safety Board. Safety Study: Evaluation of U.S. Department of Transportation Efforts in the 1990s to Address Operator Fatigue. Technical Report Safety Report NTSB/SR-99/01 May 1999 PB99-917002 Notation 7155, NTSB, Washington, DC United States of America, 1999. <http://www.nts.gov/Publictn/1999/SR9901.pdf>.
43. NHS Expert Group on Learning from Adverse Events in the NHS. An organisation with a memory. Technical report, National Health Service, London, United Kingdom, 2000. www.doh.gov.uk/orgmemreport/index.htm.
44. C. Perrow. *Normal Accidents: Living with High-Risk Technologies*. Princeton University Press, Princeton, NJ, United States of America, 1999.
45. J. Reason. *Managing the Risks of Organizational Accidents*. Ashgate Publishing, Aldershot, UK, 1997.
46. S.D. Sagan. *The Limits of Safety: Organisations, Accidents and Nuclear Weapons*. Princeton University Press, Princeton, NJ, United States of America, 1993.
47. A. Scrivener. Special report on ladbroke grove: 'pass the signal - pass the blame'. *The Locomotive Journal*, pages 8–9, June 2000. Quoted extracts from evidence to Lord Cullen's inquiry into the Ladbroke Grove accident, http://www.aslef.org.uk/dox/loco_june_00.pdf.
48. G. Sitwell and S. Purcel. Assessment of Investigations into Signals Passed at Danger (SPADs). Technical Report BL2077 004 TR06, WS Atkins Rail Limited, under contract from the HSE, London, UK, 2001. <http://www.hse.gov.uk/railway/spad/spadrep1.pdf>.
49. P. Snowdon and C.W. Johnson. Results of a preliminary survey into the usability of accident and incident reports. In J. Noyes and M. Bransby, editors, *People in Control: An international conference on human interfaces in control rooms, cockpits and command centres*, pages 258–262, Savoy Place, London, United Kingdom, 1999. The Institute of Electrical Engineers. Bath, UK, 21-23 June 1999.
50. Systems Safety Society: New Mexico Chapter. System safety analysis handbook. Technical report, Systems Safety Society, Unionville, VA, USA, 1997.
51. T.W. van der Schaaf, D.A. Lucas, and A.R. Hale. *Near Miss Reporting as a Safety Tool*. Butterworth-Heinemann, Oxford, United Kingdom, 1991.

52. W. van Vuuren. *Organisational Failure: An Exploratory Study in the Steel Industry and the Medical Domain*. PhD thesis, Institute for Business Engineering and Technology Application, Technical University of Eindhoven, Eindhoven, The Netherlands, 2000.
53. D. Vaughan. *The Challenger Launch Decision*. Chicago University Press, Chicago, United States of America, 1996.
54. C. Vincent, S. Taylor-Adams, and N. Stanhope. Framework for analysing risk and safety in clinical medicine. *British Medical Journal*, pages 1154–1157, 1998.