Normal Accidents: A Statistical Interpretation

A N Cutler, sigma engineering partnership

Royal Statistical Society conference on Risk, June 1999, Warwick, UK

Abstract

Perrow (1984) introduced the concept of "normal accidents" arising from his sociological study of hazardous incidents in supposedly high-integrity systems. Normal accidents turn out to have a natural interpretation within the statistical theory of process capability.

Such faults within systems are often overlooked because of their very complexity. Conventional design-analysis methods offer only limited protection This paper uses ideas from the Shewhart-Deming theory of process management to propose statistical tools that can safeguard complex systems against normal accidents.

1. Introduction

Perrow (1984, p5) defined a *normal accident* or *system accident*:

The odd term *normal accident* is meant to signal that, given the system characteristics, multiple and unexpected interactions of failures are inevitable.

In Perrow's model, the two characteristic features of a normal accident are:

- Unexpected and complex interactions between faults that are tolerable individually
- *Tight coupling* allowing little opportunity for mitigation or defence once a fault occurs.

These characteristics of normal accidents are illustrated by the loss of the bulk carrier MV Derbyshire (DETR, 1998). While riding out a tropical storm, an unsecured hatch cover allowed the Derbyshire's bows slowly to flood. Inability to see the hatch cover from the bridge and difficulties in operating flood-lights robbed the crew of the opportunity to understand the situation. Remedial action would have been most hazardous in the weather conditions. The flooding was not fatal in itself but reduced the elevation of the vessel's bows exposing hatch-covers on the deck to extreme wave conditions. The ship ultimately foundered when a critical hatch cover failed under severe dynamic loading from high waves crashing onto the deck. Water entering through the damaged hatch led to sudden and catastrophic flooding. None of the individual faults (hatch cover unsecured, inability to secure hatch cover subsequently, non-conservative design of critical hatch cover) would have resulted in disaster. It is most difficult to foresee such interactions at the design stage. The speed with which the flooding occurred and the prevailing weather conditions ensured that the failure would inevitably lead to fatalities. Such is the meaning of tight-coupling between a failure and its catastrophic consequences.

Perrow contends that conventional methods of hazard analysis and design assurance are inadequate to eliminate such accidents. The very complexity of the hazardous interactions entails that they remain unsuspected. Such failures to anticipate emergent behaviours in novel technologies have been well documented since classical times by Petroski (1992).

Perrow saw the solution to the threat of normal accidents in economic terms advocating market reforms to incentivise more conservative behaviour in organisations. However, Perrow's assumption is that organisations have the know-how to reduce the risks of their operations but just choose not to. It is unlikely that any business would tolerate the threats to its reputation, markets and profitability realised by the real-life incidents that Perrow describes in his book. It is more likely that, as Perrow elsewhere asserts, they simply could not see the risk. The view that managers are amoral calculators has been challenged by Vaughan (1997, p47). Reason (1990a, p216) quotes the conclusion of Wagenaar and Groeneweg after reviewing 100 shipping accidents:

Accidents do not occur because people gamble and lose, they occur because people do not believe that the accident that is about to occur is at all possible.

In this case, the situation can only improve if we can equip organisations and risk professionals with new tools with which to manage operations more lucidly.

2. Chance- and Special-Causes

We find a remarkable parallel to Perrow's insights in Harry Alpert's coining of the term common cause for faults that are inherent in a system. W Edwards Deming (1982, p134) credits Alpert with observing, in 1947:

A riot occurs in a certain prison. Officials and sociologists turn out a detailed report about the prison, with a full explanation of why and how it happened here, ignoring the fact that the causes were common to a majority of prisons, and that the riot could have happened anywhere.

The term "common-cause fault" exposes us to the danger of some confusion with the term "common-mode fault" in the risk analysis field so we shall adopt Shewhart's original term chance cause for those faults that arise from the inherent and predictable behaviour of a stable system. The idea that there are two kinds of variation:

- Chance cause: predictable (at least approximately in frequency)
- Special cause: from outside the system and inherently unpredictable, not even probabilistically

is usually associated with Walter Shewhart (1931, 1939). It also plays an important part in the thinking of the economists Frank Knight and John Maynard Keynes and can be traced back at least as far as Gottfried Liebniz in 1703 (Bernstein 1996, p118). Shewhart and Deming (1982) used the distinction as the theoretical basis of a critique of classical statistics and the development of a radical structure for business improvement and organisational management.

The distinction between the two causes of failure is important as it elucidates a fundamental danger in the misattribution of the cause of an accident as either:

- something that an individual, in the system, did or did not do
- an inevitable consequence of the system design.

In fact, Fiske & Taylor (1984) called this the fundamental attribution error.

3. Process Capability

We can understand the importance of the distinction a little more by looking at the statistical theory of process capability. Accidents are outcomes of a system and we cannot work on outcomes. We need to understand the processes of which they are the consequences. Failure can be defined as (Bignell & Fortune, 1984, p8):

A failure is said to occur when disappointment arises as a result of an assessment of an outcome from an activity.

Now, the (somewhat euphemistic) term "disappointment" suggests some legitimate personal, market or societal expectation that the process has failed to meet. We call this the voice of the customer (or stakeholder). Shewhart and Deming insisted that effective process-management begins with a separate understanding of:

- The voice of the customer
- The voice of the process.

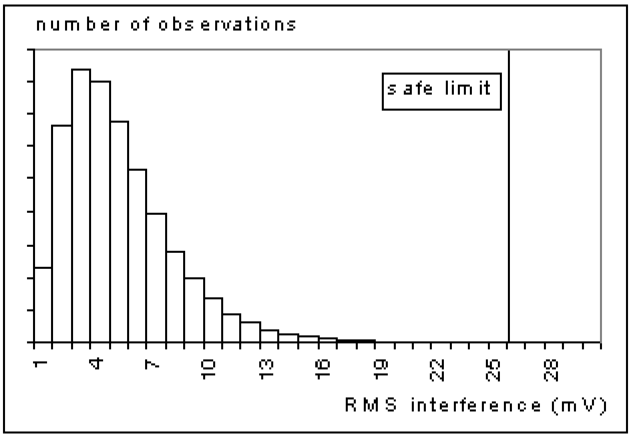

Figure 1: Histogram illustrating a capable process

This leads us to deeper insights and wider opportunities for action than simply thinking of "ok" and "accident". In particular, it leads management to the central task of continually improving the alignment between the two voices. Let's look at a simple example that illustrates a normal accident situation. Suppose that we wish to be assured that a signalling system on a railway will not operate in a hazardous manner through interference from other electrical equipment in the vicinity. The voice of the customer here will be something like "RMS potential across input terminals must not exceed 25 mV (in some specified frequency band.)"

How can we now characterise the voice of the process? Well, we could measure the percentage of the time that such an interference level is exceeded on the railway (or some other representative system). Such censored data is very weak for managing processes. We only get much data when there are many accidents. Alternatively, we could measure the actual potential and present the data as a histogram (Figure 1).

Now, the picture leads us to two questions:

- Is this picture representative of future behaviour of the system?
- What are the safety implications of the picture?

That the historical picture is representative of the future can only be assessed through our collective technical and economic understanding of the system, supported by evidence of stability from control charts. So long as we feel able to make such a judgement of stability we can make an assessment of the capability of the process. Assigning probabilities in the tail of the process variation seems hopeless owing to:

- The wholly unquantifiable behaviour of the tail beyond observed data
- The uncertainty of future process-stability.

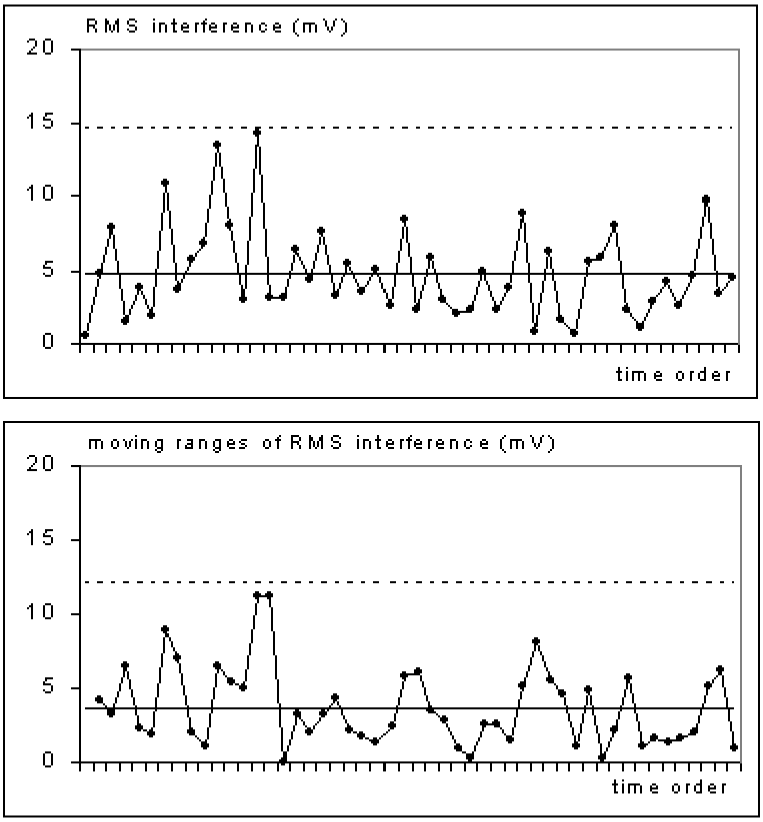

Figure 2: Control chart of RMS interference for the histogram in Fig 1

The picture tells us the range of variation in the process and compares it to the system-safety requirements. We can see what sort of variation in performance is unsurprising. The picture in Figure 1 suggests a reasonable margin of safety between the voice of the process and the voice of the customer that will protect us against:

- Chance-cause variation in the process
- Some future instability and special-cause variation outside the experience base.

However, persistent control-charting is necessary continually to validate our expectation of stability. Frequent signals of special-cause variation challenge us to review the system models that supported our earlier judgements of stability and safety. Within complex and tightly-coupled systems, process stability can only truly be assessed with the help of control charts. Control-charts (Wheeler & Chambers 1992) show the range of chance-cause variation in a system, based on historical experience. They support the assessment of stability through:

- Signalling process-behaviour outside the existing experience base
- Setting such behaviour in context allowing appraisal by craft masters.

We always need to accompany any assessment of process-capability with a control chart (Figure 2). The control limits (horizontal lines) show the process variation based on data, not our hopes or expectations. They are, therefore, sensitive to process change and instability and show signals of encroaching trouble before system integrity is violated.

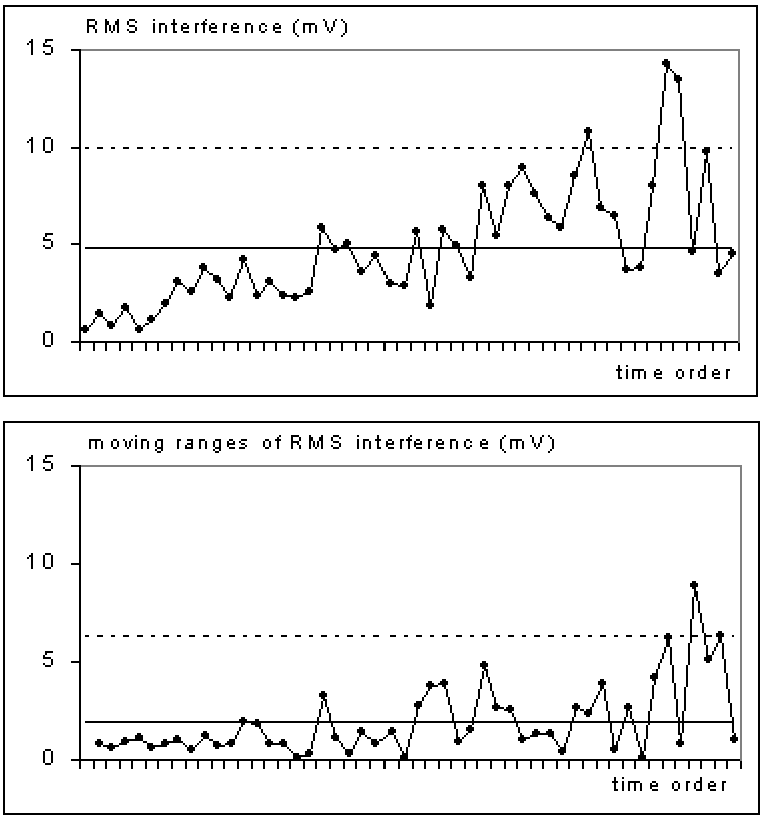

Figure 3: Control chart for a drifting process with histogram in Fig 1

The importance of the control chart is illustrated by Figure 3. The control chart in Figure 3 has exactly the same histogram (Figure 1) as does that in Figure 2. There is a general increasing trend justifying the fear that the specification may ultimately be violated.

Now consider the alternative histogram in Figure 4.

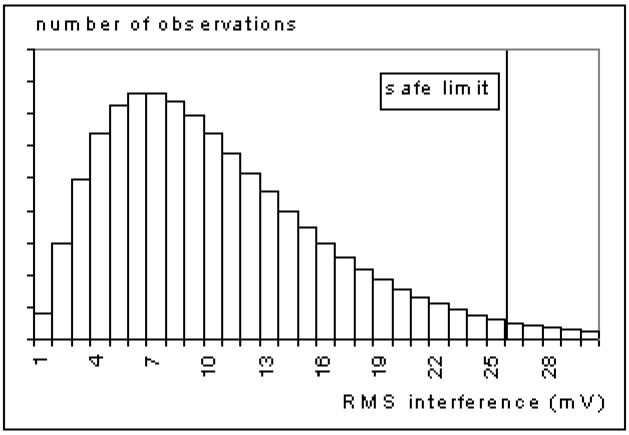

Figure 4: Histogram illustrating an incapable process

In this case, a fault in the signalling system will be disappointing but unsurprising. We don't need any special cause, any external event, any degradation, any human failure or component malfunction to cause an accident. In this sense, the accident is normal. The variation of which it represents an extreme limit is constantly present in the system.

The capability picture is a much more powerful means of assessing system integrity than is an analysis based on rate-of-occurrence-of-faults (ROCoF). Any assessment based on ROCoF can only accumulate data as quickly as accidents occur. It can only describe a system rather than support improvement and accident prevention. Understanding failure in terms of process variation is fundamental to improvement.

4. The Fundamental Attribution Error

Suppose that we were in the situation represented by Figure 4 above but without the capability picture, simply with a train wreck. Without the insight that the accident was an inevitable consequence of the system design, we would be led to ask questions like "what happened?" and "who was to blame". Now the fact is that there will be many ways in which the circumstances surrounding the accident will differ from the routine on the nonaccident days. Such happenstance data leads accident investigations to find causes and to publish recommendations and condemnations. This is the fundamental attribution error in the prison riot example above. Anyone who has been a participant in the bead-box experiment (Deming 1982, p346) will have felt this keenly.

The danger of fundamental misattribution is that much greater when we are exposed to variation in human, rather than technological, performance. Even though we may struggle to quantify human behaviour in the same straightforward way as RMS voltage, the principle of process variation and its alignment with system demands is the same. As Vaughan (1997, p393) puts it:

> ... we bring to our interpretation of public failures a wish to blame, a penchant for psychological explanations, an inability to identify the structural and cultural causes, and a need for a straightforward, simple answer that can be quickly grasped. But the answer is seldom simple. Even when our hindsight is clear … as long as we see organisational failures as the result of individual actions our strategies for control will be ineffective, and dangerously so.

Such investigations make the most elementary of statistical mistakes: testing effects suggested by the data - using the same data for exploratory and confirmatory data-analysis - in the words of Diaconis (1985) magic thinking. Reason (1990a, p212) asserts a similar error in the conventional accounts of accidents as diverse as Three Mile Island, Chernobly, Challenger, Bhopal and Zeebrugge.

The fundamental attribution error may lead us to miss root causes of failure within the system design and to ruin individuals who have displayed nothing other than normal carelessness. There is a further danger that investigations spawn recommendations founded on happenstance data and lead to tampering with the system. By tampering we mean interference with a stable system in such a way as to increase variation and hence risk. Such behaviour is well illustrated by the Nelson funnel experiment (Deming 1982, p327). As Reason (1990a, p173) puts it:

> While operators can, and frequently do, make errors in their attempts to recover from an out-of-tolerance state, many of the root causes of the emergency were usually present within the system long before these active errors were committed.

Tampering frequently occurs when we observe a disappointment or out-of-tolerance condition in a system. Sometimes, this is simply the result of chance-cause variation as illustrated in Figure 4. Unacceptable chance-cause variation is often a signal that we need additional defensive measures and, ultimately, reduction in variation through some fundamental system redesign.

We can also fall into the opposite error of viewing special-cause variation as though it arose from chance-causes. NASA engineers did not identify the association between unexpectedly low launch-pad temperatures and O-ring failures in the space shuttle booster rockets. They interpreted this critical signal as simply chance variation in the failure of the joints. Lack of this insight was critical in the decision to launch the Challenger on its final and disastrous flight (Vaughan 1997, p383).

Effective process improvement is known to demand that we make a good job of distinguishing signals of chance- and special-cause variation (Wheeler & Chambers 1992). The importance of making such a distinction is to:

- inform investigation of root causes and remedial actions
- guard against well-meaning tampering with a stable system
- advise when accusations of negligence or wrongdoing might be justified
- clarify when further defensive measures are required in the system

5. Complex Systems

Of course, any competent engineer will have spotted the importance of the capability analysis of the interference problem. Extending such thinking to human factors presents more of a challenge. In the same way that there is chance-cause variation in electromagnetic interference there is also variation in all aspects of human performance. The sort of capability pictures that we drew in Figures 1 and 4 are also models of human perception, attention and diagnostic ability, even where we have as yet no simple way of quantifying these characteristics.

Perrow argues that normal accidents commonly arise when there are unexpected interactions between factors at the extremes of their range of chance-cause variation. Thus they have the property of chance-cause failures in that nothing has happened to the system and they are inevitable consequences of the design. However, they also share something of the character of special-cause failures in that they are unpredictable given our historical knowledge base.

The range of accidents described in detail by Perrow (1984), Bignell & Fortune (1984), Reason (1990a) and Petroski (1994), among many others elsewhere, remind us of the potential for complex systems, even when subject to the most stringent design-analysis and conscientious operation, to display unanticipated and emergent behaviours. However, one is also struck by the fact that in many such cases, investigators have observed that data was already present in operations contradicting the design principles on which the safety of the system was premised.

Reason (1990b) suggests that new management techniques are necessary to guard against such latent design flaws, especially where human factors are critical. The Shewhart-Deming theory of process management offers the means by which a system may, at the earliest opportunity, detect behaviours outside the experience base and take appropriate defensive and remedial action (Cutler 1997, Cutler & Goddard 1998). Furthermore, it protects us against the tendency continually to tamper with a stable system owing to our disappointment with the outcomes and misplaced belief in our ability to control the chance-cause variation. Langer (1982) describes such a belief as the illusion of control and points out that the more complex the system, the more seductive the illusion.

Organisations sometimes balk at the commitment to measurement that such a strategy entails. However as Arrow (1971) has pointed out, the commercial value of the system knowledge that we gain by measuring a process is frequently greater than the return from the direct process outputs. Rigorous application of an appropriate strategy for variation reduction is well known to improve effectiveness and to reduce costs. It can similarly improve safety.

Once we are honest about our limited ability to predict the behaviour of complex systems from historical experience, we see the need for a programme of continual monitoring and improvement. The following steps are needed:

a) Develop the voice of the customer

This is more than some quantitative specification of the tolerable frequency of accidents. What does society count as an accident? What is the impact of varying levels of toxin in the environment? What is the impact on the economy of a shutdown? What are the public's expectations? How will the regulator judge matters? Ultimately, we need to develop the requirements into operationally defined specifications to work to on a day-to-day basis. However, this must not blind us to continual improvement towards, for example, never ending reduction in emissions, routine or accidental.

b) Develop outcome measures

When working to improve safety and profitability we will doubtless introduce many changes into the system. Unfortunately, not all of theses are guaranteed to represent improvement. We need continually to measure the overall performance of the system in order constantly to answer the question "How will we know when a change is an improvement?"

Such measurement needs to go beyond counting accidents. Resnikoff's paradox is the term often used to refer to the conflict that high-integrity systems have few faults and hence there is little data on which to base claims of integrity. Measuring physical outputs, as in the railway interference example, offers a solution to this dilemma.

c) Develop process measures

However, simply measuring outputs leaves little scope for mitigation or defence when the system exhibits surprising behaviour. We need to measure further upstream in the process so that we are monitoring, not only accidents, but also failure tokens. Such measurement:

- signals instability in upstream process measures often giving early warning of system failure allowing timely preventive or defensive measures
- supports the growth of stability and capability in upstream processes, central to improving system safety
- allows continual improvement of the critical upstream management and decision-making processes, as advocated by Reason (1990a, p210)
- characterises the unexceptional range of variation in performance arising from chance-causes.

d) Assess stability

As we observed above, within complex and tightly-coupled systems, process stability can only truly be assessed with the help of control charts. Control-charts (Wheeler & Chambers 1992) show the range of chance-cause variation in the process, based on historical experience. They support the assessment of stability through:

- Signalling process-behaviour outside the existing experience base
- Setting such behaviour in context allowing appraisal by craft masters.

Landau & Stout (1979) observe:

... in an intelligently managed organisation, the information generated by anomaly, by discrepancy between expected and actual outcomes, becomes the means by which fallible rule sets are corrected and moved toward solution sets.

Signals of special causes of variation challenge the experience base and ought to lead us to question our understanding of system behaviour. It is such signals that present the opportunity to identify emergent behaviours before they coincide with the conditions that will make an accident inevitable. The presence of frequent signals of special causes implies that the system is fundamentally unpredictable and that we are exposed to unknown and unquantifiable risks unless we introduce further defensive measures.

e) Assess capability

Just because a process is stable, it does not mean that it is acceptable. A stable system of trouble is common in many organisations. For example, the system in Figure 4 needs additional defensive measures before it can be viewed as safe.

Ultimately, we need to work on improving the system by reducing variation. By reducing variation and improving capability we can eliminate the costs of defensive measures thus achieving the goal of improved safety and reduced cost.

f) Institute continual monitoring, learning and improvement

All this measurement is of no value unless it is systematically applied to challenge design assumptions, build knowledge and trigger improvement. Within an organisation criticism of design and operating principles is a defence against confirmation bias. The quality of reporting requires cultural changes in organisations. Such changes are known to be achievable (O'Leary & Chappell 1997, Reason 1998). In particular, the authority to investigate and act needs to be passed as near to the sphere of operations as is possible. Reason (1990a, p212) describes such systems as generative organisations, capable of responding quickly to signals of danger and realising high levels of effectiveness, efficiency and safety. Elsewhere such organisations have been described as learning organisations. Perrow's original analysis assumed that such delegated authority was incompatible with the operation of a tightly-coupled system and that only command and control measures were feasible. However, Rochlin et al. (1987) cite evidence to the contrary in their analysis of US nuclear aircraft-carrier flight-deck operations.

6. Conclusions

Improvement in the integrity of complex systems demands:

a. Our acceptance that the performance of complex systems cannot be predicted at the design stage with adequate accuracy to guarantee safety over the lifetime of the asset.
b. The continual monitoring of system performance through measures that provide more insight than accident statistics.
c. That performance monitors must be put to work through organisations committed to continual improvement through learning and action, using the right statistical tools.

7. References

Arrow, K J (1971) *Essays in the Theory of Risk-Bearing*, North Holland

Bernstein, P L (1996) *Against the Gods: The Remarkable Story of Risk*, Wiley

Bignell, V & Fortune, J (1984) *Understanding Systems Failures*, Open University Press

Cutler, A N (1997) Deming's vision applied to probabilistic risk analysis, Second Edinburgh Conference on Risk: Analysis, Assessment and Management, September

Cutler, A N & Goddard, E J (1998) Aspects of "stress-strength" analysis for electricity generation and demand data, 13th Advanced Reliability Technology Symposium, IMechE/ SaRS, April, Manchester

Deming, W E (1982) *Out of the Crisis: Quality, Productivity and Competitive Position*, Cambridge University Press

Department of the Environment, Transport and the Regions (1998) *M. V. Derbyshire Surveys: UK/ EC Assessors' Report - A Summary*, HMSO

Diaconis, P (1985) Theories of data analysis: from magical thinking through classical statistics, in Hoaglin et al., (eds) *Exploring Data Tables Trends and Shapes*, Wiley

Fiske, S T & Taylor, S E (1984) *Social Cognition*, Addison-Wesley

Landau, M & Stout, R (1979) To manage is not to control: or the folly of type II errors, *Public Administration Review* **39**, 148-156

Langer, E J (1982) The illusion of control, in D Kahneman et al.(eds) *Judgement under Uncertainty: Heuristics and Biases*, Cambridge University Press, p231

O'Leary, M & Chappell, S L (1997) Confidential incident reporting systems create vital awareness of safety problems, *ICAO Journal* **51**, 11-13

Perrow, C (1984) *Normal Accidents: Living with High-Risk Technologies*, Basic Books

Petroski, H (1994) *Design Paradigms: Case Histories of Error of Judgement in Engineering*, Cambridge University Press

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

Reason, J (1990b) The contribution of latent human failures to the breakdown of complex systems, *Philosophical Transactions of the Royal Society of London* (Series B) **327**, 475-484

Reason, J (1998) Achieving a safe culture: theory and practice, *Work & Stress* **12**, 293-306

Rochlin, G I et al. (1987) The self-designing high-reliability organisation: Aircraft carrier flight operations at sea. *Naval War College Review*, Autumn

Shewhart, W A (1931) *Economic Control of Quality of Manufactured Product*, Van Nostrand

Shewhart, W A (1939) *Statistical Method from the Viewpoint of Quality Control*, Graduate School, Department of Agriculture, Washington

Vaughan, D (1997) *The Challenger Launch Decision: Risky Technology, Culture and Deviance at NASA*, Chicago University Press

Wheeler, D J & Chambers, D S (1992) *Understanding Statistical Process Control*, second edition, SPC Press