

Safety Assurance: Fact or Fiction?

Carl Sandom

iSys Integrity Limited
10 Gainsborough Drive
Sherborne, Dorset, DT9 6DR, England

carl@iSys-Integrity.com

“If the facts don't fit the theory, change the facts” Albert Einstein (attributed).

Abstract

Many safety-related systems are also socio-technical systems and providing safety assurance for these systems is extremely challenging. Providing comprehensive safety assurance evidence for the technical elements of anything but the simplest of systems is impossible due to the complexity involved and these difficulties increase dramatically when the human and organizational factors have to be considered. Apart from the inherent complexity associated with the development of safe socio-technical systems, there are other reasons to believe that safety assurance claims can be overly optimistic and based more upon fiction than fact.

This paper will examine where improvements could be made to the safety assurance process. The paper will first consider some of the reasons why safety assurance claims may be based too much upon ‘self-fulfilling prophecies’ appealing only to confirmatory and highly subjective evidence because of inherent methodological limitations with the safety assurance process and an overreliance on professional judgement. The paper will then examine a significant but common area of neglect for safety assurance claims; specifically, the widespread fixation on technology despite the prevalence of socio-technical issues for many safety-related systems. Finally, suggestions will be made regarding how to improve the validity of safety assurance claims through the use of metaevidence.

Keywords: argument, claim, evidence, induction, metaevidence, professional judgement, safety assurance, socio-technical.

1 Introduction

Systems engineering is hard enough without adding to the complexity; yet the use of socio-technical systems in high-risk environments is prevalent despite the fact that these systems often contain a complex mix of hardware, software and firmware designed, operated and maintained by people and organisations within highly-dynamic

environments often using complicated rules and procedures. The rapid rate of technological change and the use of emerging technologies in safety-related environments have also brought with it added complexity for systems engineers and new or improved processes are required to maintain the status quo.

Safety assurance is often claimed with reference to a safety argument supported by evidence that a system is acceptably safe; this broad framework for making safety assurance claims has been around for some time and is now the generally accepted paradigm within the safety engineering discipline. This paper challenges some of the fundamental assumptions underlying the current safety assurance paradigm and argues that there are some major limitations with this approach regardless of the particular safety standard or guidance adopted.

The aim of this paper is to stimulate debate on the limitations associated with safety assurance claims made for systems which are too often overly reliant upon subjective judgement and incomplete evidence to support tenuous claims regarding mainly the technical aspects of socio-technical systems safety.

Many safety assurance process improvements could be suggested; however, this paper will restrict itself to an examination of three significant and prevalent shortcomings namely: methodological limitations; professional judgement and technology fixation.

2 Methodological Limitations

Without wishing to get too deep into the philosophical discussions regarding questions of reasoning and knowledge (see Hume (1777), Popper (1959) and Kuhn (1962) for detailed discussions); it is useful for systems engineers to consider the common approaches that underpin reasoning and the acquisition of knowledge; we do this to focus on the limitations associated with the approaches used to reason about safety. (Note: there is no definitive view on the validity of knowledge; this paper will restrict itself to the prevalent view which has prevailed since the mid 20th Century. Also, some intentional simplifications are made here for the sake of brevity).

2.1 Problems of Induction

There are two broad approaches to reasoning known as deductive and inductive. Briefly, deductive reasoning

Copyright © Australian Computer Society 2011. This paper appeared at the *Australian System Safety Conference on Managing Systems and Software Safety Risks in Emerging Technologies*, Melbourne. Conferences in Research and Practice in Information Technology (CRPIT), Vol. TBD. T. Cant, Ed. Reproduction for academic, not-for profit purposes permitted provided this text is included.

progresses from the general to the specific. Deductive reasoning begins with a theory which is then refined into more specific hypotheses that can be tested. Specific hypotheses are further refined by collecting supporting observations. Finally, hypotheses are tested with specific data and the original theory is either confirmed or rejected. In contrast, inductive reasoning works the other way, moving from specific observations to broader generalizations and theories. Inductive reasoning begins with specific observations which suggest certain patterns or trends. From these patterns, tentative hypotheses (note the word tentative for the discussion later) are formulated from which general conclusions or theories are developed (Figure 1).

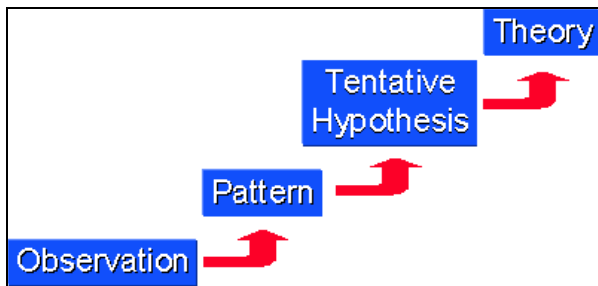


Figure 1 – Inductive Reasoning Stages

Deduction and induction processes are inextricably linked as, at some point one relies upon the other for validation. For example, a deductive safety argument may claim that a system is safe (hypothesis) then construct an argument based upon evidence (observations) to support the original claim; at some point the process will reverse and become inductive to validate the original deductive claim and vice versa.

Both inductive and deductive methods have been used for reasoning about safety even if systems engineers don't recognize those terms or use the same terminology; as discussed, the current practice for reasoning about safety assurance is for a claim (hypothesis) to be made with reference to a safety argument (pattern) supported by evidence (observation). Consequently, it is argued here that any limitations with the basic scientific approaches are also limitations with the safety assurance process.

The first significant work on the *problem of induction* was attributed to Hume (1777) and later refined by Popper (1959); Hume raised the important question of whether inductive reasoning actually does lead to knowledge and the main limitations of induction can be simplified as (Okasha 2001):

1. Hypothesizing about patterns or trends based on some number of observations can be flawed as it only takes one counter observation to nullify the hypothesis (e.g. the inference that "all swans we have seen are white, and therefore all swans are white," before the discovery of black swans). Put another way, making a safety assurance claim based upon an argument supported by some arbitrary quantity of evidence may lead to a false claim as the safety engineer may have overlooked the 'black swan' piece of evidence.

2. Past data tells you nothing about the future; therefore, it is possible that the future will turn out differently from how we believe; therefore knowledge of the future is impossible. All experimental conclusions proceed upon the basis that the future will conform to the past. Or, to put it another way, any safety assurance claim is based upon evidence that *suggests* a certain outcome based upon our past experiences; but the suggestion may be false (again, the black swan).

The problems of induction have been understood and generally accepted since 1777 but, despite that, there have been major scientific advances based upon inductive and deductive reasoning. If we accept that the problems with induction are irrefutable, and most philosophers and scientists do, we could conclude that the limitations are academic and meaningless in the context of engineering methods; however, for safety engineering at least, this is not so as flawed hypotheses may lead to unexpected failures and catastrophic accidents.

2.2 Tentative Hypotheses

For safety assurance purposes we must be proactive in trying to identify the flaws in a safety assurance claim otherwise a claim made at the outset that a system is safe may simply become a self-fulfilling prophecy as supporting evidence is sought to the exclusion of any counter-evidence that may negate the claim. Put simply, safety claims should be considered only as *tentative hypotheses* until strongly challenged by attempts to prove them false.

Kinnersly (2011) puts forward a similar opinion and suggests an alternative view to the accepted safety assurance paradigm; he argues that scientific methods should be adopted in safety engineering whereby a safety claim is examined from the view of hypothesis and challenge rather than the current norm whereby a claim that a system is safe is shown to be true as a logical consequence of appropriate (or compelling) evidence. One of the findings of the Haddon-Cave report (2009) into the loss of a UK Nimrod aircraft in Afghanistan made numerous criticisms of the way safety claims are made and concluded that the safety assurance process is not 'new' suggesting that well established (i.e. old) scientific methods have relevance for the current paradigm.

These points are consistent with the assertion made here that the inherent problems with induction should lead to a change in approach for safety engineers to challenge tentative hypotheses by proactively seeking evidence to counter claims made about the safety of a system.

2.3 Black Swans

The term 'Black Swan' is used in philosophy as a metaphor for something that hasn't been observed and therefore its existence is assumed to be improbable but not impossible. The term originates from the ancient Western conception that all swans that had been observed were white and (by the logic of induction) it was

therefore concluded that black swans could not exist. However, black swans do exist and they were first discovered in Australia in the 17th Century. Taleb (2008) takes the metaphor further and raises the prospect of 'Black Swan Events' which he characterizes as:

1. Having a central and unique attribute and high impact; his claim is that almost all consequential events in history come from the unexpected, yet humans later convince themselves that these events are explainable in hindsight.
2. The probability of rare Black Swan Events cannot be computed using scientific methods owing to the nature of the small probabilities involved.

Taleb (2008) makes the general point regarding the shortcomings of the inductive scientific method and makes a case for a new approach which attempts to answer improbable "what if" questions which he refers to as 'counterfactuals'. Interestingly, Perrow (2011) used a counterfactual approach (although he didn't refer to it in these terms) when he predicted almost exactly the failure mode of the recent Fukushima Daiichi reactors (Ladkin (2011)):

"A hurricane could take out the power, and the storm could easily render the emergency generators inoperative as well" (Perrow 2011, p134);

"No storms or floods have as yet disabled a plant's external power supply and its backup power generators". (Perrow 2011, p173).

The failure modes were evidently not foreseen by the Fukushima safety engineers as a claim was made that the Fukushima plant was acceptably safe; however, a counterfactual safety argument like Perrow's could have challenged that assertion. Clearly there is a degree of hindsight to this now, and safety engineers typically deal only with 'credible' issues but the general point being made here is that safety-related systems developers should question, justify and document what is assumed to be credible and consider potential Black Swan events.

2.4 Summary

The key point made here is that the collection and analysis of safety evidence should be based on proactively and explicitly challenging any claim that the system is safe rather than merely seeking evidence to confirm it. To paraphrase Kinnersly (2011), safety professionals need to adopt a 'challenge the claim' mentality to safety assurance rather than accept self-fulfilling arguments backed up only by confirmatory evidence. In addition, the boundaries of credibility should be challenged and Black Swan events considered; after all it is usually improbable events such as those at Fukushima that are found to be the primary causal factors for most major disasters.

It has been argued here that safety assurance evidence can be deficient due to the inherent problems of induction and improbable events; however, it is also argued that the evidence that is presented can be over-reliant on

professional judgement which is also an inductive process.

3 Professional Judgement

Professional judgement (or expert opinion) can be defined as the ability of a person or group to draw conclusions, give opinions and make interpretations based on a combination of evidence from diverse sources such as experiments, measurements, observations, knowledge and experience (McKenna and Mitchell 2006). Professional judgement is frequently used by systems developers of all disciplines and it relies upon a combination of impartial and biased facts and opinions and, for anything but simple scenarios, subjectivity can be hard to discriminate from objectivity. For example, the problems of perception when applying professional judgement to decisions on risk have been well documented (see Adams 1995).

Professional judgement is often used when an expert doesn't have any accurate or statistically significant data and the order of magnitude required for the solution to be acceptable is estimated by applying judgement gained through a combination of: academic training; experience and professional development. Professional judgement can be considered poor if highly subjective evidence is accepted as fact without consideration of where or how the evidence is derived and without an appreciation of when it is invalid. Safety assurance claims are founded upon professional judgement and it is useful to consider examples of how conclusions, opinions or interpretations may be derived from incomplete or inadequate evidence.

3.1 Statistical Inference

Safety assurance claims often need to be made for systems which are fielded before the existence of sound empirical data and claims are therefore based upon a high degree of professional judgement. In the absence of empirical data, systems developers must make statistical predictions *a priori* when, for example, considering technical or human failure rates and their associate risks. Clearly, professionals do not need to be 100% certain about something before it can be considered *a priori* knowledge; however, the point made here is that making safety claims based upon subjective judgements for which there is little evidence must be avoided; particularly in safety-related systems.

However, that is not always the case, professional judgement may be applied for example for software safety assurance and some level of inferred safety integrity may be claimed based upon evidence of software reuse in an evolving product which has been fielded on multiple platforms over a significant period of time. However, claims based upon software reuse can be based upon flawed assumptions; for example, the software (and perhaps even the hardware platform) may have been subject to considerable changes for maintenance or improvement over the period of time considered effectively invalidating any claims.

Statistical inference can lead to systems safety claims based upon a circular argument whereupon a judgment is

based on a probability when the probability was based on judgement. Vick summarizes this situation neatly with the phrase:

“...subjective probability is judgement’s quantified expression” (Vick, 2002, p393)

This situation occurs throughout the safety assurance process; particularly in those analyses based upon quantitative techniques and methods where subjective opinion is based upon subjective opinion without taking into account their source.

3.2 Assurance Gap

In addition to using judgement for statistical inferences, opinion is also often used to bridge assurance gaps. Complex systems cannot be tested exhaustively to provide definitive evidence that the required standards of safety assurance have been achieved; for example, a system would need to be tested continuously for more than 10 years, under operational conditions, with no dangerous failures and no system modifications to demonstrate that it met the IEC 61508 (2010) SIL1 target of $10E^{-6} < p_{fh} < 10E^{-5}$ (Littlewood & Strigini 1993).

Thomas (2004) points out that the lowest integrity level that current safety standards consider safety-related are associated with a probability of dangerous failure per hour that is in practice too low to be demonstrated and therefore engineering judgement must be applied by various professionals to justify claims made about systems safety. If a system cannot be exhaustively tested, the resulting assurance gap must be bridged with reference to professional judgement which, as history has shown, is not infallible.

3.3 Summary

For these reasons, and many others, safety assurance is ultimately a matter of professional judgement. Safety-related system developers in particular have a responsibility to show that where professional judgement has been applied and, for safety assurance claims, that it must be defensible. The application of professional judgement is a necessity for any systems development; however, it remains problematic; particularly for safety-related systems development.

It has been argued here that safety assurance evidence can be deficient due to methodological limitations with the safety assurance process and also that safety claims may be over-reliant on professional judgement. However, perhaps the most significant limitation for safety assurance claims is the widespread fixation on technology even for obvious socio-technical systems.

4 Technology Fixation

A socio-technical system is a system composed of technical and social sub-systems or elements; for example, Air Traffic Control Centres or Nuclear Power Stations are socio-technical systems with people organized into social structures, such as teams or departments, to do work for which they use technical sub-systems like radars, computers, radios etc. The term

‘*socio-technical system*’ and the socio-technical approach to systems design was first used by Eric Trist (1981) and presented as a radical alternative to the scientific management approach (Taylor 1911).

The socio-technical systems approach is devoted to the effective integration of both the technical and social systems and these two aspects must be considered together for safe systems development because what is optimal for one component may not be optimal for the other and design trade-offs are required. Paradoxically, the prevalent approach to safety-related systems development is often to design the technical ‘system’ and let the operators and maintainers adapt to it. It is useful to consider why safety-related system developers do not always address the socio aspects as well as the technical.

4.1 Scope & Complexity

Many safety-related systems are socio-technical systems; yet, they are often developed predominantly by systems engineers and often have little or no explicit input from human or organizational factors experts. As well as traditional systems engineering expertise, knowledge is also required from other disciplines such as human factors and organizational factors experts to ensure that socio-technical systems are designed to balance the trade-offs necessary for safe systems.

Simplistically, a socio-technical system may be considered a combination of people and technology; however, they are much more complex. Consider the typical elements that comprise a socio-technical system and the full diversity of expertise required to provide safety assurance for each element (Computing Cases 2011):

1. **Hardware and software.** These elements are likely to be an integral part of any socio-technical system. Software often incorporates social rules and organizational procedures as part of its design making them difficult to identify and to change in safety-related systems. Providing safety assurance for system hardware and software elements is relatively easy compared with the non-technical elements.
2. **People.** Individuals, groups, roles (e.g. support, training, management, engineer etc.). People can exert a positive and a negative influence on system safety and humans can alternatively be considered as ‘hazard’ or ‘hero’ depending upon the circumstances (Sandom 2007). Ideally, an interdisciplinary approach should be taken to safety-related systems development through an integrated application of Human Factors and Systems Engineering methods and techniques.
3. **Procedures.** Official and actual procedures, management models, reporting relationships, documentation requirements, rules and norms are all parts of a system and can affect its safety. Procedures describe the way things are done in an organization (or at least the official version of how they should be done) and their analyses are

essential for understanding complex socio-technical systems.

4. **Laws and regulations.** Laws and regulations are like procedures but they carry special societal sanctions if the violators are caught. Regulations are often the basis upon which system requirements are derived and they must be taken into account for the design and maintenance of the other system elements throughout the life of a system.
5. **Environment.** The complexities of the environment within which a system operates must be taken into account for any safety assurance claim. This includes aspects such as weather, and other physical conditions within which the socio-technical operates.

This vast scope, and the resulting complexity, presents a challenge for systems developers who need to consider the safety-related aspects of the entire system and then to focus the limited resources available on the most critical system functions.

The scope of any safety assurance claim must cover all these elements for socio-technical systems. If the risks associated with the non-technical elements are not considered a system will not achieve the required level of safety assurance. If the mitigations provided by the non-technical elements are not considered the technical elements may be over engineered at unnecessary cost to achieve a target level of safety assurance.

4.2 Summary

In the absence of a holistic approach to socio-technical systems safety assessment, it is tempting to concentrate safety assurance effort on what we understand or think we understand (such as hardware and software) and to adopt a 'head in the sand' approach to the human and organizational factors which are often perceived as too difficult. Humans are often the major causal factor for hazards in safety-related systems (Sandom 2002) and yet human failures often don't receive proportionate attention in safety analyses. On the other hand, human operators also often provide substantial mitigation between machine-originated hazards and their associated accidents; yet this too is often overlooked or, conversely, sometimes over-stated.

Perhaps the most significant shortcoming of many safety assurance claims is the widespread fixation on technology. The conclusion to be drawn from this is that in many instances safety claims at best provide only limited safety assurance as the prevalent errors in socio-technical systems are often related mainly to issues associated with human and organizational factors.

5 Improving Safety Assurance

From the previous discussions, it was asserted that there are some significant limitations on the veracity of the evidence supporting safety assurance claims which are caused by methodological limitations, professional

judgement and technology fixation. A safety claim can be backed up with a perfectly logical argument but still fail to provide assurance if the evidence is inadequate (McDermid 2001). The main aim of this paper is to stimulate debate on the limitations associated with safety assurance; however, some suggestions will now be made on how to improve the validity of safety assurance claims through the use of what is described here as metaevidence.

The prefix 'meta' is used to describe a concept which is an abstraction from another concept; for example metacognition could be described as 'thinking about thinking'. Assertions have been made in this paper regarding the perceived shortcomings of safety assurance claims and, specifically, their reliance on incomplete and/or unconvincing evidence. To address the shortcomings described, it is suggested here that metaevidence (i.e. evidence about evidence) should be sought to support a claim that safety assurance evidence is both comprehensive and compelling.

5.1.1 Comprehensive Evidence

Some general improvements can be made to the safety assurance process by ensuring that the scope of the safety evidence is comprehensive by addressing the issues previously discussed. Specifically, metaevidence should be sought to take into consideration the following:

1. **Challenge Claims.** Evidence should be actively sought to challenging systems safety claims rather than simply focusing upon confirmatory evidence which is the norm. A review of three of the major safety standards in common use today revealed that only UK Defence Standard 00-56 (MoD 2007) contains a requirement to consider counter-evidence and this is not developed further in the guidance (Kinnersly 2011). Pragmatically, this will require a sufficient degree of independence in the overall safety assurance process as the person(s) responsible for making a safety claim are not well placed to try breaking a safety claim; the same principle is applied for independent validation and verification in systems engineering.
2. **Consider Black Swan Events.** Safety assessments must necessarily be bounded and it is normal practice to focus only on what is perceived to be credible; however, the bounds of credibility need to be agreed and evidence should be presented to back up all related assumptions made by systems developers. Something that may be considered incredible during system development may be considered probable later in the operational life of the system so assumptions must be revisited periodically in light of emerging technologies and other changes. Analysing the incredible may seem like an unnecessary task; however, a brief examination of many disasters will reveal that the improbable has actually occurred (e.g. Fukushima).

3. **Examine Subjectivity.** All safety assurance activities rely upon professional judgement which is inherently subjective and should therefore be critically examined to ensure that the resulting safety claims are reasonable and remain so over time. Statistical inference is particularly sensitive to error for quantitative analyses (e.g. Fault Tree or Human Reliability Analyses) and the assurance gap created by a lack of testing is another area of focus. Systems developers should seek evidence that any professionals applying professional judgement to safety assurance claims are competent to do so. In addition, it is equally important to ensure that the application of professional judgement to safety-related issues is not simply the opinion of a single person and a consensus from a group of competent professionals should be formed.
4. **Extend Scope of Analyses.** The scope of systems safety assurance activities should be extended from the norm to include all elements of socio-technical systems which requires expertise and contributions from different disciplines (e.g. engineering, sociology, cognitive psychology etc.). Ignoring the non-technical aspects of many safety-related systems has a significant impact on the actual safety assurance provided. Programme managers should ensure that interdisciplinary teams are formed for the analysis of safety in socio-technical systems; despite the lack of regulation or guidance in this area provided by the primary safety standards. Consider the simple reality that in some domains human factors account for more than 90% of accident or incident causal factors (Sandom 2004); yet the human factors are often not been properly addressed making system safety assurance claims fictional.
2. **Insufficient data.** A common problem with evidence sampling is drawing conclusions from insufficient data; this is related to the problem of induction (see 2.1). It is not enough to observe a couple of instances of data that support a safety claim; however, it is not easy to decide how much data is statistically sufficient. Sufficiency of data is a matter of degree; the more evidence the better and the amount of confidence that we can have in an inference grows gradually as more evidence is brought in to support it.
3. **Unrepresentative data.** Simply having a lot of data is not enough to guarantee that a claim is valid; it is generally important that the data has been drawn from a representative sample of sources and obtained under a variety of different conditions. For example, it may not be enough to show that requirements-based testing has been undertaken for software, a valid claim may also require some proof of absence of errors during operation of the system. Special attention should be paid to evidence relating to evolving products where claims are made based on past performance without properly considering the impact of configuration changes or changes in the context of use. For example, software safety evidence taken from use in fixed wing aircraft may not be valid in rotary winged aircraft.

5.1.2 Compelling Evidence

In addition to questions of comprehensiveness, safety evidence should be assessed to determine if it is convincing. The credibility of safety evidence should be assessed to determine where it comes from and if it is adequately representative of the claims being made. Metaevidence should be sought to take into consideration the following possible evidential criteria:

1. **Misrepresenting data.** Data can be deliberately or unintentionally represented. Data can be misrepresented deliberately by claiming that it suggests something when it does not; this can be the case with safety evidence for example when programmes are under severe pressure to meet budgets, milestones and targets. A further way in which data may be misrepresented is if it is presented selectively and a varied data set is described by focusing only on certain sections of it. Data can be unintentionally misrepresented as conclusions are hurriedly based upon initial evidence found to fit a given proposition.

5.1.3 Summary

In summary, it is suggested that metaevidence should be sought to support a claim that safety assurance evidence is both comprehensive and compelling before a system is operational and throughout the operational life of a system.

It is recognised that metaevidence is itself evidence and it can be argued recursively that it should also be comprehensive and compelling and require evidence to demonstrate that it is so. However, at some point the law of diminishing returns must apply and professional judgement (or consensus opinion) must be applied to bring the process to a halt when little value is being added. Nonetheless, it is asserted here that at least one-level of metaevidence should be sought for all but the simplest safety-related systems.

6 Conclusions

There are many safety-related, socio-technical systems in operational use today and many of these have based safety assurance claims on inductive arguments, a great deal of professional judgement and have only considered the technology; yet, thankfully there are few catastrophic accidents or serious incidents associated with these systems. The relatively small number of catastrophic accidents or serious incidents associated with these systems could lead us to conclude that our safety assurance processes are sufficiently robust; however, this is not the case.

A safety claim will usually be made relative to an acceptable level of risk and it is suggested here that a

great deal of uncertainty and sensitivity of these claims can be attributed to the issues raised in this paper. A safety claim is *not an incontrovertible fact* and the nature of the safety assurance process means that it is often difficult to determine the robustness or validity of a claim. It is often impossible to determine how close to being unsafe a system might be.

From the arguments presented in this paper, it may be concluded that it is not possible to provide valid system safety assurance without major professional input from sociologists and cognitive psychologists and without using sound scientific methods. However, safety professionals shouldn't 'throw the baby out with the bathwater' as, despite the issues raised in this paper, there are relatively few accidents given the vast number of complex, safety-related systems in existence.

Although there is room for improvement in current safety assurance best practice it is not suggested here that a paradigm shift is required, merely an evolution of the existing practice to address the major limitations, some of which have been discussed in this paper, and to enable safety professionals to better separate fact from fiction.

7 References

- Adams, J (1995): *Risk*. Routledge, London.
- Computing Cases: <http://computingcases.org>. Accessed 26 April 2011.
- Haddon-Cave, C. (2009): *An independent review into the broader issues surrounding the loss of the RAF Nimrod MR2 aircraft XV230 in Afghanistan in 2006*. Her Majesty's Stationery Office.
- Hume, D. (1777): *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*, Niddich, P. H. (Ed.) 3rd Edition 1975. Oxford University Press.
- Kuhn, T, S. (1962): *The structure of scientific revolutions*. University of Chicago Press.
- IEC 61508 (2010): *Functional safety of electrical/electronic/programmable electronic safety related systems*. International Electrotechnical Committee. Ed. 2, 2010-04. Geneva, Switzerland.
- Kinnersly, S. (2011): Safety Cases – what can we learn from Science? in Dale, C., and Anderson, T. (eds), *Advances in Systems Safety, Proc. Safety-Critical Systems Club Symposium 2011*, Springer-Verlag, London.
- Ladkin, P. (2011): posted on Safety Critical Mailing List, <http://www.cs.york.ac.uk/hise/safety-critical-archive/2011/>. Accessed 26 April 2011.
- Littlewood, B. and Strigini, L. (1993): Validation of Ultra-High Dependability for Software-based Systems in *Communications of the ACM* **36** (11) 69-80.
- McDermid, J. (2001): Software Safety: Where's the evidence? *6th Australian Workshop on Industrial Experience with Safety Critical Systems and Software (SCS'01)*, Brisbane. Conferences in Research and Practice in Information Technology, Vol. 3 P Lindsay, Ed.
- McKenna, S. and Mitchell, J. (2006): *Professional Judgment in Vocational Education and Training: A Set of Resources*. 2nd Ed. Commonwealth of Australia, Department of Education, Science and Training.
- MoD (2007): Defence Standard 00-56 Issue 4. Safety Management Requirements for Defence Systems: Part 1 Requirements; Part 2 Guidance on Establishing a Means of Complying with Part 1. UK Ministry of Defence.
- Okasha, Samir. (2001): What did Hume really show about induction? *The Philosophical Quarterly*, **51** (204).
- Perrow, C. (2011): *The Next Catastrophe: Reducing Our Vulnerabilities to Natural, Industrial, and Terrorist Disasters*. Princeton University Press.
- Popper, K. R. (1959): *The Logic of Scientific Discovery*, New York: Basic Books.
- Sandom, C. (2002): Human Factors Considerations for System Safety, in *Components of System Safety*, Redmill F and Anderson T (Eds.), proceedings of 10th Safety Critical Systems Symposium, 5th-7th February 2002 Southampton, Springer-Verlag, UK.
- Sandom, C., and Harvey, R. S. (2004): *Human Factors for Engineers*, The Institution of Electrical Engineers, UK.
- Sandom, C. (2007): Success and Failure: Human as Hero – Human as Hazard. Conferences in Research and Practice in Information Technology (CRPIT), Vol. 57. T. Cant, (Ed.), *12th Australian Conference on Safety Related Programmable Systems*, Adelaide.
- Taleb, N. N. (2008): *The Black Swan: The Impact of the Highly Improbable*. Penguin.
- Taylor, Frederick Winslow (1911), *The Principles of Scientific Management*, New York, NY, USA and London, UK: Harper & Brothers
- Thomas, M. (2004): Engineering Judgement. Conferences in Research and Practice in Information Technology, Vol. 38. Cant, T. (Ed), *9th Australian Workshop on Safety Related Programmable Systems*, Brisbane.
- Trist, E. L. (1981). *The evolution of socio-technical systems: A conceptual framework and an action research program*. Ontario Quality of Working Life Center, Occasional Paper No. 2.
- Vick, S. G. (2002): *Degrees of Belief: Subjective Probability and Engineering Judgement*, American Society of Civil Engineers Press.