How did systems get so safe without adequate analysis methods?

J A McDermid*, A J Rae †

**University of York, York, UK john.mcdermid@york.ac.uk, † Griffith University, Brisbane, Australia*

Keywords: Weaknesses in safety analysis; good safety records; high reliability organisations; complexity.

Abstract

We are remarkably safe. Across a wide range of industries fatal accident rates have been falling for decades, and accidents are less significant as a cause of death than at any time during the 20th century. This is despite the fact that we are using analysis techniques which were developed decades ago for applications that were far simpler than the complex technical and socio-technical systems we produce today. In fact, the situation is rather worse than this. Examination of industrial practice shows that these inadequate techniques are often even more inadequately applied. These concerns apply, *mutatis mutandis*, to safety standards. This leads us to enquire "how did systems get so safe?" If we can answer the question (or shed light on it) then this should help us to focus effort on methods that are effective, and thus to preserve safety in the face of growing complexity of, and societal dependence on, safety-related systems.

1 Introduction

Society is highly dependent on safety-related systems, the majority of which are now computer controlled. Many of these systems are already complex, and the complexity is only growing as systems become more highly interconnected (the terms cyber-physical systems (CPS) and Internet of Things (IoT) are often used). However, despite this complexity, in major areas of activity these systems are remarkably safe. Two particular examples are civil air transport and railways.

The good safety record of systems on which society depends is, whilst desirable, actually rather surprising, as the methods we use for safety analysis are limited, and their application is often weak. The limitations of the "tried and trusted" methods have driven research into better approaches, e.g. [1]–[3], but these advances have had fairly little impact on industrial practice. Put another way, industry is using techniques that the leading researchers believe are not fit for purpose. This seems, to us, to be a paradox and leads to the question: "How did systems get so safe without adequate analysis methods?"

In our view some of the answers are analogous with what has been found in high reliability organisations (HROs) that were characterised as having a "better than expected" safety record [4]. In this paper we take the premise that many organisations have a "better than expected" safety record, if the adequacy of safety analysis activities alone sets the benchmark.

To substantiate these points we first set out, in section 2, some evidence for the claim about the level of safety achieved by modern systems and, in section 3, address the issue of the growing societal dependence on such systems.

Section 4 presents a brief analysis of the perceived problems with current methods and their industrial application. This cannot be comprehensive but aims to identify sufficient issues that it should convince the reader at least to be somewhat more circumspect about what they expect from standards and analysis methods.

The core of the paper is section 5. It sets out illustrations of those issues we perceive as contributing to achievement of safety. This is supported by evidence, as much as possible, but many of the observations reflect experience we have had either acting as consultants or gleaned through teaching and are not supported by references. In some cases it is easier to explain what makes systems safe by giving counter-examples where those contributing factors were absent; in these cases it is not possible to identify companies or projects.

We take the HRO perspective, particularly in section 5, with two purposes in mind. First, to build a hypothesis about the primary causal factors in achieving safety which might be more explicitly tested through empirical evaluation. Second, to identify the types of risk that are likely to arise as the complexity of our systems grows. These issues are drawn out in section 6. Section 7 presents our conclusions.

2 Safety Achievements

To back up our opening claim that "we are remarkably safe" we consider two specific domains. Some might argue that other domains could have been chosen where there are less exemplary safety records. That is true, but life expectancy in the developed world continues to grow to the extent that it is viewed as a "demographic crisis" [5]; at minimum we can say that safety-related systems aren't reducing life expectancy.

2.1 Aerospace

In civil aerospace, Boeing has long collected and published aircraft accident data across the worldwide fleet; [6] is the latest annual data. This shows a trend from the early days of flying, through the introduction of modern jets to the latest generations (with substantial in-service lives), to a very low accident rate of less than one per million departures which has been sustained for several decades, despite growing aircraft complexity and growing traffic volumes.

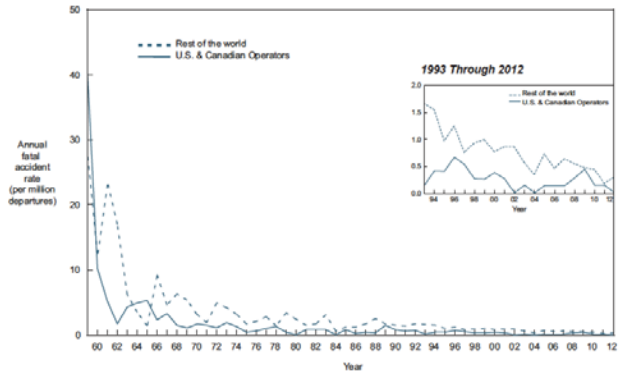

Figure 1: Aircraft Accident Rate Trends (from [7])

Despite these good overall trends there have been problems for the industry, for example the early losses of aircraft of the Airbus A320 family, e.g. [7], and recent incidents with the Boeing 787 [8]. These inform our analysis of the reasons for high levels of safety achievement.

2.2 Railways

In the rail sector, similar high levels of safety performance are seen, although it is hard to find "global" data. However, the UK Railway Safety Standards Board (RSSB) publishes data giving some view of the European situation.

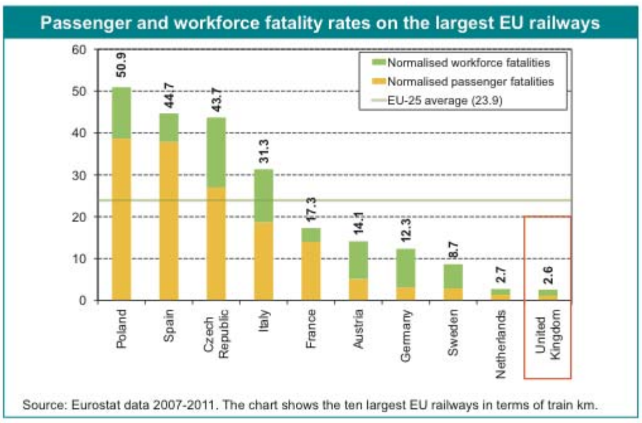

Figure 2: UK Data from RSSB (from [9])

This shows that accident rates can be very low, but the impact of local factors, which might be training or track state, can be seen in the spread of fatality rates.

3 Societal Dependency

It is clear that society is dependent on safety-related systems; but there is merit in outlining the scale of our dependency and the growing reliance on computer-based systems, e.g.:

- Healthcare – implanted devices, assisted living for the elderly, and hospital services, e.g. diagnostics, etc.;
- Logistics – delivery of food to supermarkets, fuel to service stations, etc.;
- Transportation – buses, cars, trains, aircraft;
- Manufacturing – most sectors including pharmaceuticals have highly automated production;
- Utilities – electrical power, gas, water (controlling production/generation and distribution) and sewage.

There is growing complexity, and arguably also dependency, in many of these sectors. Three illustrations will suffice.

In healthcare, there is work on systems for assisted living, see for example [10], to enable society to deal with the demographic trends mentioned above. The aim of these technologies is to enable the elderly (and often infirm) to continue to live independently – in their own homes – rather than needing residential care (be it state or privately funded).

In the automotive sector, the desire for increased autonomy is apparent. This is not just the Google car; major car companies have projects of varying degrees of ambition, and Gothenburg is planning to allow limited autonomous driving in 2015 [11].

In utilities, there is a growing trend to "smart networks" with much more complex control over consumption and a move in some countries, e.g. Germany, to a model of a "prosumer" – producer-consumer – who both takes power from the grid and sells it back to the grid [12].

The latter two, at least, should be viewed as being systems of systems (SoS) where safety depends on the interaction between systems, as well as the functioning of each system.

4 Limitations

It would be a massive undertaking, far beyond the scope of this paper, to do a thorough review of the limitations of safety analysis methods and standards. Instead the aim is, at least, to give the reader reasons to be cautious about accepting claims about the efficacy of current methods and standards.

The authors have analysed use of quantitative risk assessment (QRA) and shown key weaknesses [13]. There is no evidence to support a claim that the figures for (unsafe) failure rates are accurate; weaker hypotheses, e.g. that quantitative methods are better than qualitative ones, are not well supported either. It would be desirable to present such analyses for other methods, but to the authors' knowledge no such results exist, thus the remainder of this section gives a rather more judgmental treatment of the issues but hopefully gives enough information to act as a "cautionary tale".

4.1 Methods and Standards

Classical methods, e.g. failure modes, effects and criticality analyses (FMECAs) [14] and fault trees [15] consider failures of components and systems, using hierarchy to manage scale. However the signature of modern incidents and accidents is one of interactive complexity, not component failures. At best these methods help with parts of an analysis. Even fairly simple systems, have failure behaviours involving temporal dependencies which such methods cannot address [16].

More fundamentally, even single integrated circuits are too complex to conduct FMECA, and the bulk failure rate data used does not really apply to modern designs with reduced feature sizes, voltage levels, etc. Without a solid basis, the system-level analyses are questionable.

Standards make (often implicit) assumptions. IEC 61508 [17] has a model of systems with an equipment under control and a protection system; yet many modern systems are highly integrated; what does it mean to apply such a standard to an integrated control system without separate protection?

ISO 26262 [18] in the automotive sector uses controllability as a key factor in determining safety integrity levels; how applicable is this for a car when the driver has no controls?

Also, there is growing concern with cyber-security as a "threat" to safety (this is a consequence of the use of CPS and the IoT) but is not widely addressed by standards.

4.2 Engineering Practice

Weaknesses of methods and standards are only a small part of the problem. Although both authors have seen good industrial practices, there are numerous examples that fail to derive even the available benefit from the methods and standards.

A "fault tree" for a communication-based railway signalling system had a top event "system fails" and this was broken down to sub-system level, again just in terms of "failures". This was developed before the design detail was produced but was not updated as the design progressed. Although there was no treatment of hazards in the fault tree (some were recorded in a hazard log) this was accepted by independent assessors!

In aerospace, some analyses (based on fault trees) produced probabilities of exactly $1x10^{-9}$ for all the catastrophic hazards – the figure required by the standards. If the failure rate distribution is normal then there is a 50% chance the target isn't met (the $1x10^{-9}$ is a point probability representing the mean of a distribution) but what was more worrying was that all the calculations came out to this figure; truly precision engineering – or numerology.

Some safety cases are "tick-box" documents – for example showing "compliance" with the requirements of EN 50128 [19] by saying that the required methods for the SIL had been used, but saying nothing about hazards, safety requirements, treatment of test failures, etc.

4.3 Complexity and SoS

As indicated above, many modern accidents are more due to interactive complexity, than individual failures; the collision near Überlingen [20] is another example of this. The methods and standards in force do not deal well with such situations, hence the research initiatives identified earlier, e.g. work on how to validate the safety of sense and avoid algorithms, on which safety of unmanned air systems will depend [3].

5 Safe Engineering Organisations

The work on High Reliability Organisation (HRO) focuses on operations. We focus here on engineering, but include aspects of operations as, with modern systems, it is not reasonable to assume that all problems are "engineered out" and feedback from operations is needed to inform design engineering.

It would be good to identify "safety success strategies"; we seek to do this, but note that HRO theorists such as La Porte [4] refer to "evidence that contradicts" as a feature that distinguishes safe from unsafe organisations. Thus we also give "counter-examples" to help substantiate the relevance of the "positives" which we adduce.

We consider engineering, engineering management and operations in turn. In each case there is a value judgement as to the factors that are chosen; one of the aims of doing this is to construct hypotheses that are worthy of further study; we turn to this perspective in section 6.

5.1 Engineering

With regard to engineering, we identify six factors that we believe have a positive influence on safety.

1) Domain knowledge: there is evidence [21] that problems found late in development or in service predominantly arise from requirements errors; domain knowledge is a key factor in identifying and removing such errors.
    a) Counter-example: Boeing outsourced aspects of the 787 development to companies without aerospace experience, with attendant problems [8].
2) Evolving products: it is hard to identify hazards, and potential hazard causes with novel designs or technology so evolving successful designs is a way of controlling the risk due to these sources of uncertainty.
    a) Counter-example: the use of carbon-fibre for aero engine fan blades by Rolls-Royce [22] (arguably this was more of a commercial problem, but would have been unsafe if the engines had entered service).
3) System architecture: the more critical the system, the greater the architectural defences, especially redundancy, to prevent, detect and mitigate faults.
    a) A review of aerospace products found that the level and depth of defences reflected criticality, in a way which was independent of manufacturer, and which was beyond what is required by standards [23].
    b) Counter-example: the high loss-rate of drones has many factors, e.g.: "some common drone models were designed without backup safety features" [24].
4) Safety and systems engineering: safety being an effective tool in systems engineering, guiding the design.
    a) Some aircraft have had a preliminary system safety assessment (PSSA) on competing system designs to aid in selecting the most effective designs.

5) Conservatism: avoidance of “leading edge” technologies to avoid the uncertainties of novel designs.
    a) Some aircraft system suppliers use their own processor designs to avoid the complexity of modern commercial processor chips;
    b) Counter-example: carbon-fibre fan blades also can be cited here [22].
6) Control of engineering maturity: this includes design for manufacture, and prioritisation of problems based on their safety criticality, leading to removal of the more critical design defects before entry into service.
    a) Counter-example: “value engineering” replacing a part with benign failure modes with a cheaper one which had highly undesirable failure modes (without manufacturing consulting the designers).

These factors are not all completely independent, e.g. 5 and 6 are clearly related. It is also not clear that these are the only factors that are important; but there is evidence to support these points, even if it is only experiential, in some cases.

5.2 Management

With regard to management, organisation, governance and culture, we identify five factors to consider.

1) Priority of engineering: supporting a culture of “doing the right thing” almost regardless of the constraints (fix it first, worry who pays second).
    a) A “bug” was found in a flight control system where an integrator was initialised later than it should have been; although blending of control laws meant that this was only like “noise” at the control surfaces the “bug” was fixed despite the very high cost of re-verification (it was a simple and cheap change).
2) Good leadership: giving appropriate priority to safety, in terms of resourcing, process improvement, listening to concerns of engineers, etc.
    a) Counter-example: a manager who when told that he could reduce the error and rework rate by improving review checklists said “there’s no point, when we are under time pressure we don’t do any reviews”; so he had time to “do it over again” but not to do it right.
3) Supportive/just culture: listening to the concerns of engineers and the way in which the company responds to problems, rather than blaming or ignoring them.
    a) Some companies have good mechanisms for dealing with engineer concerns; an aircraft manufacturer cited an example of an employee who produced over 40 concerns in a few years, most of which were non-issues, but the 22nd was serious and needed attention so it repaid the effort of dealing with all the others.
    b) Counter-example: several accidents, with Challenger [25] being one of the more obvious.
4) Developing and rewarding competence: ensuring that individuals with appropriate skills are recruited and managed to enable them to progress, personally, and to develop the specialist expertise that the company needs.
    a) Engineering Fellowship schemes used to provide career paths for technical specialists, rather than the only route to senior levels being via management.
    b) Counter-example: a company which was short of specialist engineering skills (including safety) and was unable to recruit specialists, so it paid all current engineers for increased weekly hours on the basis that “engineers are all the same/interchangeable”.
5) Impact of regulation: and influence of knowing that products are independently evaluated, which helps to “keep the organisation honest”.
    a) The good examples cited above (aerospace and railways) are regulated, and also involve active engagement of “independent” assessors throughout the development, encouraging continuous attention to safety, not a culture of “compliance at the end”.
    b) Counter-example: there has been a relatively recent move by the US Federal Drug Administration (FDA) to require safety cases, in part due to high (unsafe) failure rates of devices.

As with the engineering factors, it is not clear that these are all the issues of concern; but there is evidence to support the relevance of these issues, even if little is in the literature.

5.3 Operations

With regard to operations, we consider five factors which have a bearing on system design; it can be argued that some of these are more “pure operational” factors but, in our view, there is a strong link to (good) design and engineering.

1) Good operators including good training: operators know how to deal with the system in normal and failure modes so that they do not initiate problems, and effectively deal with technical failures.
    a) It is common for operators of certain classes of plant, e.g. nuclear power stations, to carry out periodic emergency drills either on the main plant or on a back up system (note good engineering is required to enable this, e.g. to allow training on a live plant without endangering operations).
    b) Counter-example: the level of aircraft automation has been identified as a factor in the difficulties that pilots have found in dealing with emergencies [26].
2) Empowerment: the operators have the authority to make difficult decisions, including suspending operations (note that there is a technical perspective, as the operators need the information and the controls to be able to do this, but it is also a cultural issue).
    a) Counter-example: the accident at Wenzhou was, in part, due to pressure on operators so that they tried to continue operating through the effects of lightning and equipment design faults [27].

3) Time between initiating events and accidents: this gives operators time to assess the situation and to plan and implement appropriate remedial action.
    a) This is visible in the design of alarms and warnings philosophies of many systems.
4) Simple mitigations: despite system complexity, many hazardous failures are simple in their causes or effects, or have simple remedial actions.
    a) Counter-example: the simple response to a throttle sticking is to apply the brakes and, if necessary, to turn off the ignition; if reports are to be believed [28], use of CPS makes these mitigations infeasible (effectively creating a single point of failure).
5) Learning from experience (LfE): reporting operational issues, and removing/mitigating both the immediate causes of the issues, and doing a root cause analysis to identify and hence resolve underlying causes.
    a) Counter-example: due to pressure in a business the individual charged with managing the LfE process was moved to another post, and his was left vacant, so LfE was stopped for more than half a year during which time further operational incidents occurred.

Perhaps one of the most prevalent weaknesses, in the authors' experience, is the limitation in LfE. There are often systemic problems both in initial reporting, then in analysing the results to make informed decisions about risk.

6 Analysis

The information presented here is only partially supported by evidence but, as indicated above, the intent is to use this to form some hypotheses that can be tested. It should be noted, however, that generally there is not a single "mechanism" that ensures safety, so defining "tests" has to be done with this in mind. We begin by considering a general approach to ways in which we might study the question: "How did systems get so safe without adequate analysis methods?".

6.1 Ethnographic Studies

Ethnography is a long established, systematic way of studying people and cultures based on observation or engagement to try to understand how societies actually work, rather than how they are typically depicted. The approach has been applied to engineering; in principle it might be used to determine, for example, whether or not PSSA was critical in selecting good designs in aerospace, or the actual decision making process was influenced much more by other factors, e.g. cost, or the legacy designs (consider the point about evolving designs).

This might be a viable approach, and give value in terms of an in-depth understanding of particular industrial settings. However it would be difficult to get an understanding of the answer to the above question, as it would focus on a single organisation, the process would be slow (the studies often last many months if not years), and there may also be a difficulty of obtaining access to the industrial setting.

6.2 Analysis of Factors

An alternative is to seek to identify hypotheses from the above observations, or other sources of knowledge about industrial practices, and to frame them in ways that can be "tested" either from public-domain information, or based on focused studies or reviews that are more practical than an ethnographic study.

Not all of the points identified are easily tested because of the multiplicity of factors involved in achieving safety, but some are relatively easy to explore. For example:

Hypothesis: LfE Prevents Accidents: data from LfE systems on leading indicators of accidents is used to inform short-term remediation (e.g. changes in operating procedures) and long-term rectification of problems.

Analysis of LfE records associated with organisations or industries with low accident rates could be used to test this hypothesis; if data were available it could also be contrasted with those organisations with poor records.

Of course this hypothesis could be refuted; it might be found that industries with low accident rates also have low incident rates and that system safety is a consequence of other factors, e.g. good design and/or good operators.

There are other hypotheses that are more interesting, but harder to test, for example:

Hypothesis: Regulation is more significant than standards: the interventions made by assessors and regulators lead to safety improvements, but the interventions are not confined to adherence to the applicable standards.

This could be assessed by considering minutes of meetings with regulators, audit reports, etc. Again the hypothesis could be refuted if it were found that recommendations beyond the standards didn't exist, were never implemented, led to unused design features, etc.

Some of the engineering elements would be the hardest to test. For example, if we constructed the hypothesis:

Hypothesis: Bespoke processor designs promote safety: failure rates of complex electronic systems are lower where designers control the processor design rather than use mass-market processors. (This is given as a hypothesis, although it does not follow from the observations above, as at least two aerospace companies do design their own processors.)

This involves comparisons between alternative designs, and information on low-level causes of problems (which may not be apparent through LfE systems). It may not be easy to get such information, and there are many other factors – for example the use of modern multi-core processors is likely to support more complex functionality, and this could cloud the cause-effect relationships.

7 Conclusions

It is important to know "how systems got to be so safe". Very substantial effort is devoted to system design and analysis, writing standards, demonstrating conformance with standards, training, and so on. Regardless of whether or not the adduced important contributory factors in section 5 are important, it is clear than not all the methods, processes, etc. will be equally effective; thus we would do well to focus our efforts.

It is hard to answer the question posed in the title, but by considering how we might answer it we should begin to shed light on the issues. More significantly if we can begin to test the hypotheses then we might be able to modify behaviours for the better; if the LfE hypothesis is right, then the processes supporting LfE should be a priority for resourcing, and the counterexample mentioned would be less likely to recur. It might also be that greater benefit would accrue from a shift from "standards compliance" to regulation where independent assessors had much more engagement in development. In our view it is important to seek to answer the above question, so we continue to have remarkably good safety records, despite the growing complexity of systems on which society depends.

References

[1] E. Hollnagel, D. D. Woods, and N. Leveson, Eds., *Resilience Engineering: Concepts and Precepts*, New edition. Ashgate Publishing Limited, 2006.

[2] N. G. Leveson, 'The need for new paradigms in safety engineering', in *Safety-Critical Systems: Problems, Process and Practice*, Springer, 2009, pp. 3–20.

[3] X. Zou, R. Alexander, J. A. McDermid, 'Safety Validation of Sense and Avoid Algorithms using Simulation and Evolutionary Search', in *Proc Safecomp 2014,* Florence, Italy, 2014.

[4] T. R. La Porte, 'High Reliability Organizations: Unlikely, Demanding and At Risk.', *J. Contingencies Crisis Manag.*, vol. 4, no. 2, p. 60, Jun. 1996.

[5] European Union Center of North Carolina, 'The EU's Demographic Crisis', The EU's demographic crisis, see: http://europe.unc.edu/wp-content/uploads/2013/08/Brief9-0803-demographic-crisis.pdf.

[6] 'Statistical Summary of Commercial Jet Airplane Accidents', Boeing Commercial Airplanes, Aug. 2013.

[7] Air Accidents and Incidents Investigation Commission, 'Accident occurred on 21 May 1998 to Aircraft Airbus A-320-212 Registration GU-UKLL at Ibiza Airport, Balearic Islands', Ministry of Development, A-19/98, 1999.

[8] Boeing 787-8 Critical Systems Review Team, 'Boeing 787-8 Design, Certification, and Manufacturing Systems Review', Federal Aviation Administration, Mar. 2014.

[9] A. Clinton, 'Annual Safety Performance Report 2013/14', Rail Safety and Standards Board, 2014.

[10] J. Barlow, S Bayer, T.C. Oliviera, 'Remote Care: Health at Home', in *The Connected Home: The Future of Domestic Life*, R. Harper (ed.), Springer Verlag, 2011.

[11] see: http://www.worldcarfans.com/114042974245/volvo-launches-project-drive-me-100-autonomous-cars-hit.

[12] Siemens AG, Deliverable D6 'Future Scenarios for Smart Microgrids', version 1.0, EIT ICT Activity 11831 TSES, 2011.

[13] A. Rae, R. Alexander, J. A. McDermid, 'Fixing the cracks in the crystal ball: A maturity model for quantitative risk assessment', *Reliab. Eng. Syst. Saf.*, vol. 125, pp. 67–81, May 2014.

[14] MIL-P-1629, 'Procedures for Performing a Failure Mode, Effects and Criticality Analysis', U.S. Department of Defense, 1949.

[15] W. E. Vesely, F. F. Goldberg, N. H. Roberts, and D. F. Haasl, 'Fault Tree Handbook (NUREG-0492)', 1981.

[16] Derailment of a tram at East Croydon 17 February 2012. Rail Accident Report 130328_R042013_East_Croydon, 2013.

[17] International Electrotechnical Commission, 'Functional safety of electrical/electronic/programmable electronic safety related systems'. IEC 61508, Issue 2, 2009.

[18] International Organisation for Standardisation, ISO 26262, 'Road Vehicles – Functional Safety' – Parts 1–9, April 2011.

[19] CENELEC, EN50128, 'Railway Applications: Communications, Signalling and Processing Systems. Software for Railway Control and Protection Systems'. 2001.

[20] Bundesstelle für Flugunfalluntersuchung (BFU: German Federal Bureau of Aircraft Accidents Investigation), Accident on 1 July 2002, Near Überlingen/Lake Constance, Germany Involving Boeing B757-200 and Tupolev TU154M, Investigation Report AX001-1-2/02, May 2004).

[21] R. R. Lutz, 'Analyzing Software Requirements Errors in Safety-Critical, Embedded Systems', The Journal of Systems and Software, 1996.

[22] see: http://en.wikipedia.org/wiki/Rolls-Royce_RB211.

[23] D. W. Reinhardt, 'Safety Assurance of Aviation Systems', PhD thesis, University of York, 2013 (see: Appendix A)

[24] http://www.washingtonpost.com/sf/investigative/. 2014/06/20/when-drones-fall-from-the-sky/?hpid=z1.

[25] D. Vaughan, 'The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA', University of Chicago Press, 1997.

[26] P. J. Sherman, R. L. Helmreich, A. C. Meritt, 'National Culture and Flgihtdeck Automation: Results ofa Multi-Nation Survey', International Journal of Aviation Psychology, Vol. 7, No. 4, 1997.

[27] The investigation report on the "7.23" Yongwen line railway accident (2011) (in Chinese), The state investigation team of the China Yongwen railway accident. http://www.chinasafety.gov.cn/newpage/Contents/Channel_5498/2011/1228/160577/content_160577.htm

[28] see: http://www.suddenacceleration.com/.

All web references last accessed 22nd July 2014.