Open questions and closed minds: mapping the gaps and divisions in the safety body of knowledge

Andrew Rae
Safety Science Innovation Lab
Griffith University
170 Kessels Rd, Nathan 4111, Queensland
d.rae@griffith.edu.au

Abstract

What do we really know about safety? The volume of critique directed at current safety practice, and the academic disagreements about epistemology and research methods suggest a lack of consensus about what safety is, how safety can be investigated, and how safety can be achieved. Further, recruitment advertisements, education programs and accreditation schemes show no clear industry position on the level or type of education a safety practitioner needs. This paper suggests that divisions between safety practice and research communities can be explained by different answers to the primary research question "Why do accidents happen?" The paper explores how these answers manifest as distinct schools of thought which use a common language to conceal fundamental differences in the understanding and practice of safety.

Keywords: Safety education, Epistemology, Safety Terminology

1 Introduction

Central to every academic discipline are one or more big questions. Chemistry asks "What is matter?" Physics asks "Why does the natural world behave as it does?" Safety Science also has such a question: "Why do accidents happen?" Unlike physics and chemistry, however, new answers to the primary question of safety science seldom fully replace older answers. Whilst the solid ball atomic model of Democritus gave way to the "Plum Pudding" model of Thomson, then the nucleus model of Rutherford, the orbital model of Bohr, the shell model of Langmuir and then a sequence of progressively detailed quantum-mechanical models, new ideas in safety do not fully displace existing answers to the central question. Safety Science has experienced neither the slow march of "normal science" nor the "paradigm shifts" described by Kuhn in his "Structure of Scientific Revolutions" (Kuhn, 1996).

This is not to say that progress in safety has halted. Rather, whilst many subfields of safety have developed through the refinement of models, and replacing discredited ideas with new understanding, progress on the primary question of "Why do accidents happen?" has been in the form of adding new answers without fully discarding the old ones.

My central thesis in this paper is that communities of academic safety thought and industrial safety practice cluster around different answers to the "Why do accidents happen?" question. Since innovators tend to justify their work by critiquing the old-guard, new theories about accident causation reject the relevance of existing theories without fully discrediting them. The new ideas acquire followers at the expense of the old, but even the oldest ideas in safety have modern champions.

The different safety communities use similar language, and often adopting and adapting ideas from each other. However, as ideas cross the boundaries of safety paradigms they become transformed, sometimes subtly but occasionally to the point where they are used in direct opposition to their original intentions. Practitioners are usually trained within a single paradigm, and so may be unaware that there are substantial disagreements regarding the theory and practice of safety.

Concepts which are embraced across the safety world, but with marked differences in understanding between safety communities include:

- Safety culture
- Justice and accountability
- Risk assessment
- Systems thinking
- Safety management systems; and
- Safety Measurement

2 Disagreement without Contradiction

In this section I describe eight broad answers to the question "Why Do Accidents Happen?", and link them to academic and practice communities. Whilst there may be some overlap between the answers, and occasionally some academics and practitioners may embrace more than one answer in their work, my contention is that the different answers represent distinct schools of thought. The demarcation between the paradigms can be observed through different constructs for measuring safety, different safety practices, and intellectual communities with strong in-group identity.

None of the answers to the question are definitively wrong. Each correctly identifies a cause of accidents, in the sense that modifying the cause changes the likelihood of the effect. The schools disagree about the appropriateness and effectiveness of focussing on the particular causes.

2.1 Accidents are caused by Poor Work Habits

The "Behaviour Based Safety" approach, as exemplified by Scott Geller (2005) considers that work habits are responsible for accidents. These habits are intrinsic to individuals, and may be modified by changing the activators and consequences of specific behaviours. If individuals are conditioned to expect positive outcomes from safe behaviours, and negative outcomes from unsafe behaviours, they will follow prompts which direct them to the safe behaviours.

Behaviours can be observed (or even counted) before and after interventions, so each behavioural intervention can be tested as a workplace experiment to determine if it successfully improves individual behaviours.

The focus on individual behaviours is a development on the early idea of "accident-prone" individuals, as exemplified by Alexander (1949). This idea takes the observation that a disproportionate number of accidents occur to a sub-group of exposed populations, and concludes that individual personal attributes must be the explanation. Behaviour-based safety adds the contribution of Skinner (1965) that these personal attributes are malleable behaviours, shaped by conditioning forces that can be deliberately controlled.

2.2 Accidents are caused by Poor Collective Attitude to Safety

Another school of thought, sometimes confusingly also labelled as "Behaviour Based Safety", ascribes accidents to organisational culture. This work is exemplified by Zohar (2010), who introduced the practice of safety climate measurement (Zohar, 1980).

This school also considers that work habits are responsible for accidents, but considers these behaviours to arise from shared beliefs and values rather than individual conditioning.

2.3 Accidents are caused by System Designs that do not Adequately Control Hazards

"Safety engineering", as described by Möller and Hansson (2008), views accidents as arising from uncontrolled risk. This risk is attached to hazards – particular events, states or entities of concern – which must be managed. The focus on safety engineering is to identify, characterise, and manage the hazards through the design process. Safety engineering is a broad school of thought, since there is little consensus about the correct way to manage hazards (McDermid and Rae, 2012). Approaches include:

- Building "inherent safety" into systems by reducing dangerous forces and materials (Khan and Amyotte, 2002)
- Constructing and testing a "safety case" – a body of argument and evidence for the safety of the system (Kelly, 1998)
- Applying quantitative risk assessment to design systems with acceptable levels of risk (Aven, 2011)

These approaches share the underlying assumption that accidents arise from flawed designs, and that adoption of an appropriate design process will prevent accidents.

2.4 Accidents are caused by Failed Physical Components

The "Layers of Protection" or "Barrier" school of safety describes accidents as occurring due to inadequate protection against known harmful events. An accident begins with an "initiating event" which may be internal to the system (e.g. a process deviation) or external to the system (e.g. a lightning strike). Under normal circumstances, the disturbance will be accommodated and corrected. Failure of these protective mechanisms calls into play further protection, and so on until the disturbance is contained (Baram, 2010). Accidents are explained by inadequate barriers, an insufficient number of barriers, or by common causes that defeat multiple barriers (Hollnagel, 2008).

2.5 Accidents are caused by Inadequate Management Systems

"Safety Management Systems" (SMS) extend safety engineering to consider design and operational systems rather than the technical systems they produce and manage. Accidents are caused by inadequate control and feedback mechanisms within these secondary systems.

Leveson's Systems Theoretic Accident Model and Process (STAMP) (Leveson, 2004) is an example of this approach, although SMS concepts precede the STAMP model.

2.6 Accidents are caused by Organisational Behaviour

The organisational school of accident causation starts with the observation that accidents are failures of intelligence – viewed with hindsight, there was a period before the accident in which it could have been recognised and prevented. Major theories in this school include Turner's "Man Made Disasters" (Turner, 1976), Reason's "Vulnerable System Syndrome" (Reason et al., 2001) and "High Reliability Organisations" (La Porte, 1996).

These theories consider that disasters can be understood and prevented by examining how organisations make sense of information (Weick, 1993) and reach decisions (Cohen et al., 1972).

The organisational school of thought has been increasingly drawn to study the gap between work-as-imagined and work-as-conducted, leading to "Safety-II" approaches for safety management (Hollnagel, 2014). Safety-II suggests that accidents should not be seen as exceptional circumstances, but as one possible outcome from normal work. The focus of safety activities should therefore be on understanding and improving normal work rather than on preventing accidents.

2.7 Accidents are caused by Accepting Unacceptable Risk

It was William Lowrance who first suggested that "A thing is safe if its risks are judged to be acceptable" (Lowrance, 1976). Lowrance argued that safety was a value judgement, varying over time and between contexts. Accidents are manifestations of differences between risk judgements. This school of thought embraces the work of Fischhoff (1984) and Slovic

(2001), studying socially constructed attitudes to risk rather than means to control risk.

The theory of risk homeostasis (Wilde, 1982), whilst not universally embraced (O'Neill and Williams, 1998) translates this school into action, suggesting the management of risk by manipulation of risk perception and acceptance.

2.8 Accidents are an Inevitable Consequence of Certain Technology

The most pessimistic school of thought suggests that safety is an unattainable social goal. This view was popularised by Charles Perrow in his book "Normal Accidents" (Perrow, 1999). Perrow links accidents to two causes – "tight coupling" and "interactive complexity" – which, according to his use of the terms, are inherent properties of particular technologies. These properties prevent sufficient operational understanding to successfully manage the safety of industries using those technologies.

Amalberti (2001) also suggests that safety is unattainable, but for a different reason. He observes that as industries reduce their accident rate, the acceptability of accidents in those industries also decreases. Further attempts to improve safety may paradoxically increase the chance of an unacceptable event by raising expectations of safety that cannot be met.

3 Separated by a Common Language

In this section I discuss several concept labels in common usage, and show how they refer to different concepts when applied by the different schools of thought.

3.1 Safety Culture

"Safety culture", the shared values and attitudes relevant for safety within an organisation, has an influence on safety. Behind this basic consensus is a wide variety of views on what culture is, and the mechanism by which it affects safety.

For Behavior-Based Safety, there is a direct path from culture to behaviour to outcomes. For the "work habits" school, culture is about the conditions, motivations and rewards that drive individual actions. It is created and adjusted by deliberate acts of management, and exists within the workforce. For the "collective attitude" school, culture is sustained by the workforce through a system of norms and values. It can be transformed over time by sustained championing of new values, but is not readily malleable.

For the "safety engineering and "barrier" schools, culture is a mediating force that drives the effectiveness of safety lifecycle activities. It is the difference between activities performed according to templates, standards and rules, and activities conducted diligently and competently, with due regard for their safety impact. Culture is partly set by the expectations and reward structures of management, partly by education and understanding, and is partly sustained by the values and norms of the workforce.

For the "Safety Management System" school, there is no distinct line between culture and process. A strong safety culture is one with competent people with appropriate authority administering properly functioning safety systems. The existence of these things is, by itself, evidence of the safety culture.

In the organisational school, safety culture is closely linked to organisational learning. Safety culture exists as a sort of collective mind, where a strong safety culture is good at collecting information about itself, and a weak safety culture is not self-aware.

For the "acceptable risk" and "inevitable consequence" schools, culture is simply an average of the individual attitudes to risk. It is subject to group effects, such as risk normalisation and amplification, and may be observed in collective decision making, but it exists within individuals.

Despite these underlying differences, when asked to describe safety culture, all of these groups are likely to give similar summaries. If pressed for a definition, they are most likely to say that culture is "the way things are done around here". These superficial similarities belie fundamental differences. Each school requires different measures of culture. For each school, different strategies for influencing safety culture will be preferred, and will be observed through different effects.

3.2 Risk Assessment

"Risk assessment" is the discipline of estimating the likelihood and consequence of undesirable events. The concept of likelihood in particular suffers from apparent clarity and actual confusion.

For the "work habits" school, likelihood is measured directly by frequency. Practitioners of behaviour-based safety seldom experience uncertainty about their own knowledge of safe and unsafe behaviours. The dependent variable is simply the rate of these behaviours. A representative example is wearing personal protective equipment (PPE). Risk is expressed as the proportion of observations in which correct PPE is not worn.

For safety engineering, risk is a design-time prediction of future system properties. This prediction is formed by the imaginative anticipation of a set of hazards, and the detailed estimation of the likelihood and consequence of each hazard. Fault Trees are the stereotypical safety engineering risk assessment. The "barrier" school performs risk assessments with similar form, but with a different understanding of the source of risk. Risk is the likelihood and consequence of the failure of protective mechanisms.

In safety management systems, risk assessment is a process to be performed on prescribed occasions. Rather than a means to predict risk, risk assessment is itself a way of controlling risk. SMS risk assessments frequently use forms, templates and checklists.

The organisational school has an evolving attitude towards risk assessment. Turner (1976) viewed risk assessment as an organisational mindfulness exercise – an antidote to cultural blindness. Later authors such as Vaughan (1997) and Rae (2014) recognise risk assessment as a social exercise that can normalise attitudes to specific risks rather than warning of danger.

This position is even more strongly held by the "acceptable risk" school, which views risk assessment entirely as a social construction of meaning, rather than a way of objectively examining reality.

Unlike safety culture, the differences between notions of risk assessment are readily apparent. The various schools are more likely to accuse each other of “not doing risk assessment properly” or “not understanding risk assessment” rather than holding a mistaken belief that there is a shared concept.

3.3 A Systems Approach

“Systems approach”, “scientific approach” and “systematic approach” are the badges that almost every new safety method uses to distinguish itself from existing methods. The labels are all supposed to convey some type of merit, but the desirability and nature of a “systems approach” differs from school to school.

Behaviour based safety envisions fixed systems of work, within which humans are unreliable components. A positivist empirical approach is used to trial methods of improving human performance (Geller, 2005).

In safety engineering, the focus is on engineered systems. Hence the alternate label for the school, “system safety”. A “systems approach” to safety engineering is sometimes used to refer directly the systems engineering, but also to indicate an expanded view of what comprises “the system”, including environmental, social and managerial concerns within the system under consideration.

“System safety engineering” is not to be confused with “safety system engineering”, which forms part of the barrier approach to safety. A “systems” approach in the barrier school involves enforcement of constraints through technological rather than procedural barriers.

In the “safety management systems” school, a “systems approach” requires considering control and feedback mechanisms, rather than simply putting in place constraints. A “system” isn’t assumed to behave correctly – it has mechanisms for detecting and correcting deviations from optimal behaviour.

A “systems approach” is less unequivocally a good thing in the “organisational” school. Mindfulness and the disciplined examination of assumptions and evidence is encouraged, but a “systems approach” also carries connotations of bureaucracy. The organisational school has also embraced complexity science, suggesting that safety emerges from system interactions which cannot be decomposed or otherwise readily analysed.

The acceptable risk school recognises broader social systems with positive and negative feedback, but places less emphasis on understanding smaller technical systems.

The idea of a systems approach to safety has acquired buzzword status, and lost much of its meaning even within each school. Whilst it is generally recognised as a positive, progressive term, it can refer to contradictory approaches to safety.

4 What is a Safety Manager?

The following person description is an amalgam of posts on seek.com.au during early 2015 for the position of “Safety Manager”.

Responsibilities

- Provide guidance on workplace health and safety legislation
- Design and improve safety systems
- Plan and implement safety initiatives
- Identify and assess risks
- Respond to incidents and accidents
- Promote safety culture
- Build relationships with key stakeholders

Competencies

- Understanding of safety performance measurement
- Understanding of relevant legislation
- Written and spoken communication skills
- Analytical skills

Stripped of industry context, there is little indication whether applicants for this job should be safety engineers, behavioural safety specialists, workplace health and safety advisors, organisational psychologists, or safety systems experts. In fact, the responsibilities and competencies listed here match different posts asking for experience in each of those capacities. Each post additionally asks for experience in a relevant industry, and “suitable” (usually unspecified) tertiary education.

For someone educated or experienced in each of the different schools described in Section 2 and Section 3, the only unambiguous responsibilities or competencies are those that refer to legislation. All of the other details call for vastly different expertise depending on the approach to safety.

“Design and improve safety systems” could mean:

- Establish a practice of behaviour observation and recording
- Implement a set of rewards and punishments
- Produce a safety case for a product or service
- Engineer electronic monitoring and alarm devices
- Write procedures for hazard identification, risk assessment etc.
- Establish a practice of routine information gathering about what works well and what doesn’t work well

“Identify and assess risks” could mean:

- Observe work being conducted, and point out safety problems
- Perform statistical analysis of behavioural observations
- Apply and analyse a safety culture survey
- Lead a hazard identification workshop for a new system design
- Physically inspect plant equipment for safety problems
- Examine company procedures and point out shortfalls
- Study the organisation’s approach to understanding risk, and advise senior management on how it could go wrong
- Conduct surveys and interviews to understand how the workforce and customers perceive and relate to risk

5 The Education Solution?

One way to resolve the ambiguity identified in the previous section is via education and experience. It is possible that despite the various academic schools of thought, students are provided with a common training providing consensus and shared understanding.

There is a very limited amount of research on the education of safety professionals, mostly dating from the 1990s when there was a rapid growth in the number of postgraduate safety programs (Arezes and Swuste, 2012). However, as Arezes and Swuste note, there is a large degree of diversity amongst these programs, with little consensus about the topics to be covered.

It is likely that there is even greater diversity than is immediately apparent, since topics are typically described using terms such as “risk assessment”, “human factors”, and “safety management” which, as described in Section 4, can conceal broad differences in content and approach.

One way to investigate this further may be to consider the publication record of those teaching each program. This would not be reliable, since a researcher could publish within a particular community, but still present a range of ideas and approaches fairly. Another approach would be to compare the program reading lists. These are usually not publicly available, and further investigation was beyond the scope of this paper.

Without this detailed information, there is a high level of ambiguity concealed in program descriptions. Are students attending different universities receiving a common grounding in safety, or being adopted into closed schools of thought? Despite the use of similar terms to describe the curriculum, the output of those who go on to postgraduate study would strongly suggest training in only one school of safety thought.

6 Implications

In 2012, Nancy Leveson published “Engineering a Safer World” (Leveson, 2011), a manifesto for a new approach to safety. In 2014 Eric Hollnagel published “Safety-I and Safety-II” (Hollnagel, 2014), calling for a paradigm shift in safety thinking, and Sidney Dekker published “Safety Differently” (Dekker, 2014) ushering in a “new era” of thinking about safety. The problem with all of these revolutions is that there is no status quo to overturn.

Instead of an established paradigm, ready to be advanced or overthrown, safety science has clustered in entrenched schools of thought. Empirical investigation is slowly advancing understanding within each school, but appears incapable of settling debates between schools.

From a purely academic perspective, there is no problem to be solved. The answer to the question “Why do accidents happen?” does not need to be either simple or definitive, and can embrace all of the proposals. Practitioners, however, are frequently trained within a single school, but are then exposed, defenceless, to critiques from the other schools. How is a working professional to react when faced with multiple best-selling authors, with decades of experience in both research and polemic? The only options appear to be fight, flight or surrender – adopt an entrenched position despite all arguments to the contrary, refuse to read or listen to new safety research, or switch camps and become an enthusiastic disciple of a different school.

7 References

Alexander, F., 1949. The Accident-Prone Individual. Public Health Rep. 64, 357–362.

Amalberti, R., 2001. The paradoxes of almost totally safe transportation systems. Saf. Sci. 37, 109–126.

Arezes, P.M., Swuste, P., 2012. Occupational Health and Safety post-graduation courses in Europe: A general overview. Saf. Sci. 50, 433–442. doi:10.1016/j.ssci.2011.10.003

Aven, T., 2011. Quantitative Risk Assessment: The Scientific Platform. Cambridge University Press, Cambridge; New York.

Baram, M., 2010. Preventing Accidents in Offshore Oil and Gas Operations: The US Approach and Some Contrasting Features of the Norwegian Approach.

Cohen, M.D., March, J.G., Olsen, J.P., 1972. A Garbage Can Model of Organizational Choice. Adm. Sci. Q. 17, 1–25. doi:10.2307/2392088

Dekker, S., 2014. Safety Differently: Human Factors for a New Era, Second Edition, 2 edition. ed. CRC Press, Boca Raton.

Fischhoff, B., Lichtenstein, S., Slovic, P., Derby, S.L., Keeney, R., 1984. Acceptable Risk. Cambridge University Press.

Geller, E.S., 2005. Behavior-Based Safety and Occupational Risk Management. Behav. Modif. 29, 539–561. doi:10.1177/0145445504273287

Hollnagel, E., 2014. Safety-I and Safety-II, New edition edition. ed. Ashgate, Farnham, Surrey, UK England ; Burlington, VT, USA.

Hollnagel, E., 2008. Risk + barriers = safety? Saf. Sci., Occupational Accident Scenarios, and Accident Analysis Papers selected from the third international conference Working on Safety (WOS2006), September 12-15th, 2006, Zeewolde, The Netherland 46, 221–229. doi:10.1016/j.ssci.2007.06.028

Kelly, T., 1998. Arguing Safety - A Systematic Approach to Managing Safety Cases. University of York.

Khan, F.I., Amyotte, P.R., 2002. Inherent safety in offshore oil and gas activities: a review of the present status and future directions. J. Loss Prev. Process Ind. 15, 279–289. doi:10.1016/S0950-4230(02)00009-8

Kuhn, T.S., 1996. The Structure of Scientific Revolutions, New ed of 3 Revised ed edition. ed. University of Chicago Press, Chicago, IL.

La Porte, T.R., 1996. High Reliability Organizations: Unlikely, Demanding and At Risk. J. Contingencies Crisis Manag. 4, 60.

Leveson, N., 2011. Engineering a Safer World. MIT Press.

Leveson, N., 2004. A New Accident Model for Engineering Safer Systems. Saf. Sci. 42, 237–270. doi:10.1016/S0925-7535(03)00047-X

Lowrance, W.W., 1976. Of Acceptable Risk: Science and the Determination of Safety. Kaufmann, William, Incorporated, Los Altos, Calif.

McDermid, J.A., Rae, A.J., 2012. Goal Based Safety Standards: Promises and Pitfalls. Presented at the Safety-critical Systems Symposium.

Möller, N., Hansson, S.O., 2008. Principles of engineering safety: Risk and uncertainty reduction. Reliab. Eng. Syst. Saf. 93, 798–805. doi:10.1016/j.ress.2007.03.031

O'Neill, B., Williams, A., 1998. Risk homeostasis hypothesis: a rebuttal. Inj. Prev. 4, 92–93.

Perrow, C., 1999. Normal Accidents: Living with High-Risk Technologies. Princeton University Press.

Rae, A.J., McDermid, J.A., Alexander, R.D., Nicholson, M., 2014. Probative Blindness: How Safety Activity can fail to Update Beliefs about Safety. Presented at the System Safety and Cyber Security, IET, Manchester.

Reason, J.T., Carthey, J., de Leval, M.R., 2001. Diagnosing "vulnerable system syndrome": an essential prerequisite to effective risk management. Qual. Health Care QHC 10 Suppl 2, ii21–25.

Skinner, B.F., 1965. Science and Human Behaviour, New impression edition. ed. The Free Press, New York, NY.

Slovic, P., 2001. The risk game. J. Hazard. Mater. 86, 17–24. doi:10.1016/S0304-3894(01)00248-5

Turner, B.A., 1976. The Organizational and Interorganizational Development of Disasters. Adm. Sci. Q. 21, 378–397. doi:10.2307/2391850

Vaughan, D., 1997. The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA, 1 edition. ed. University Of Chicago Press, Chicago.

Weick, K.E., 1993. The Collapse of Sensemaking in Organizations: The Mann Gulch Disaster. Adm. Sci. Q. 38, 628–652. doi:10.2307/2393339

Wilde, G.J.S., 1982. The Theory of Risk Homeostasis: Implications for Safety and Health. Risk Anal. 2, 209–225. doi:10.1111/j.1539-6924.1982.tb01384.x

Zohar, D., 2010. Thirty years of safety climate research: Reflections and future directions. Accid. Anal. Prev., Safety Climate: New Developments in Conceptualization, Theory, and Research 42, 1517–1522. doi:10.1016/j.aap.2009.12.019

Zohar, D., 1980. Safety climate in industrial organizations: theoretical and applied implications. J. Appl. Psychol. 65, 96.