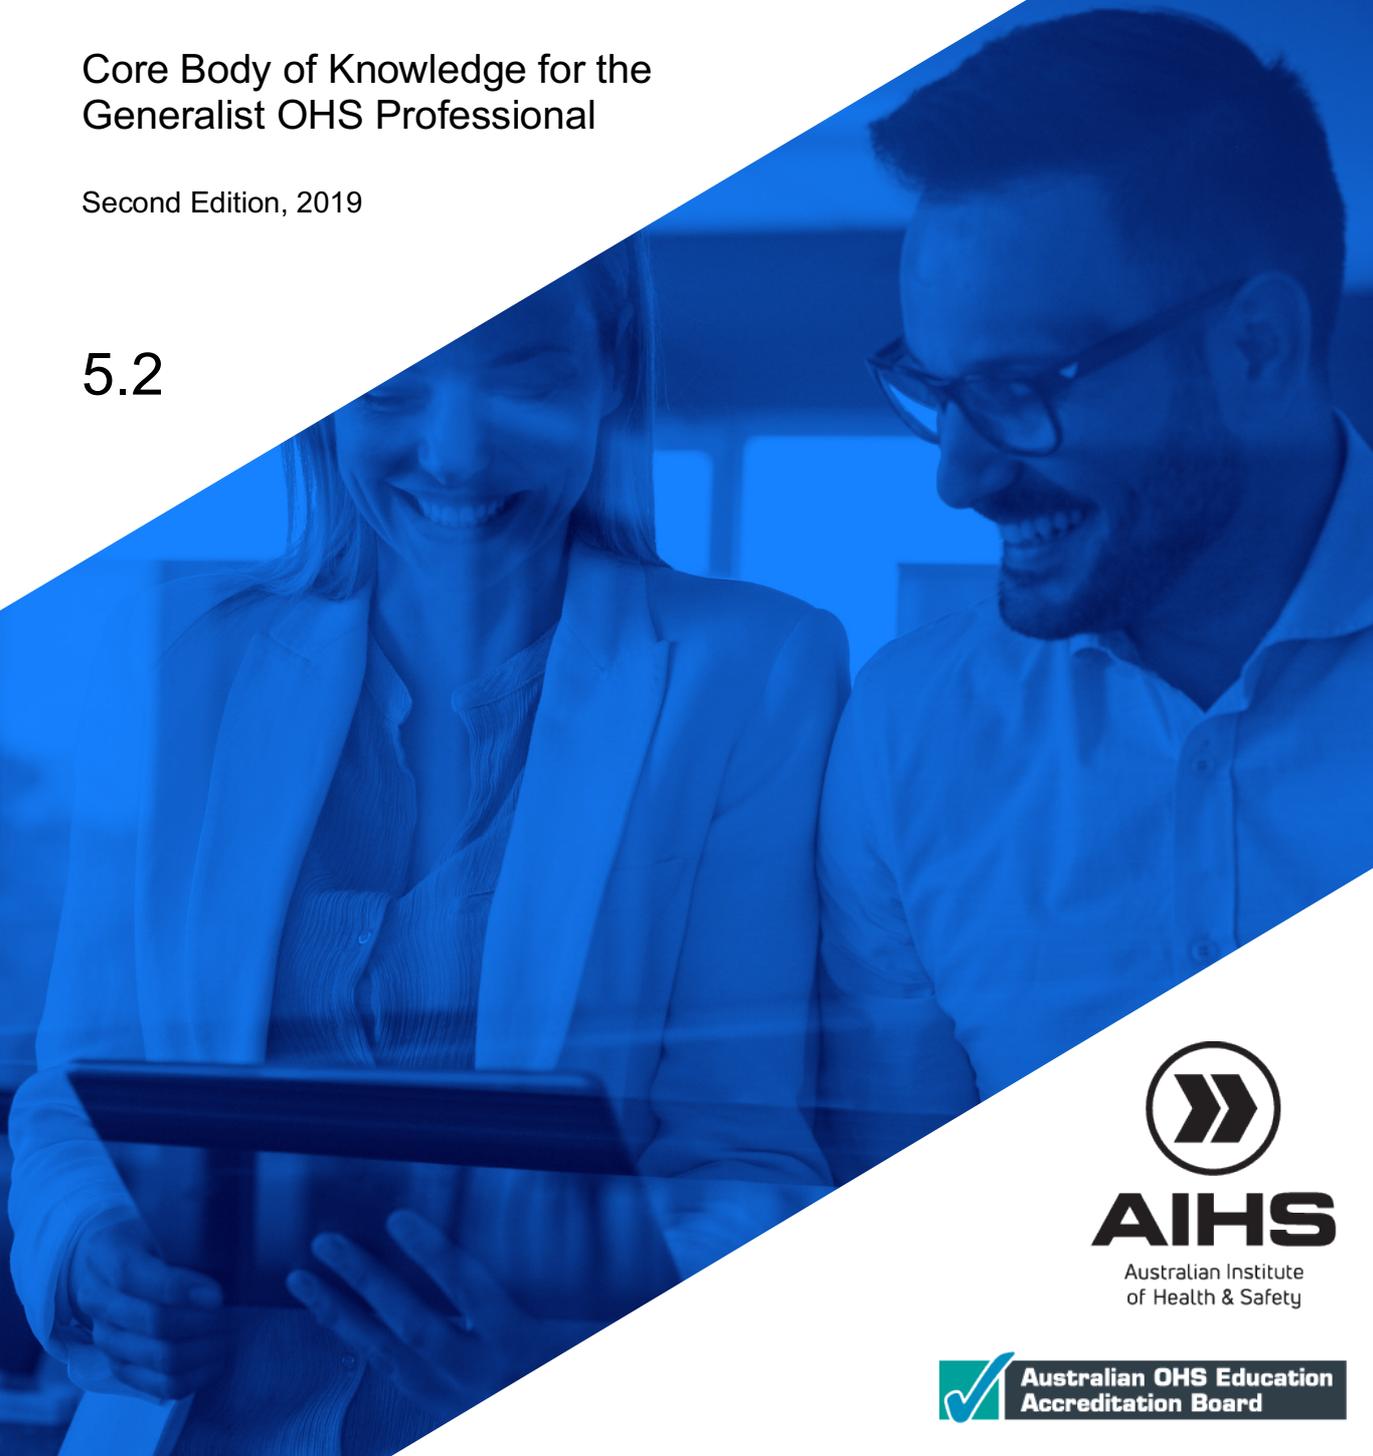


Global Concept: The Science of Safety

Core Body of Knowledge for the
Generalist OHS Professional

Second Edition, 2019

5.2



AIHS

Australian Institute
of Health & Safety



**Australian OHS Education
Accreditation Board**

Copyright notice and licence terms

Copyright (2019) Australian Institute of Health and Safety (AIHS), Tullamarine, Victoria, Australia

This work is copyright and has been published by the Australian Institute of Health and Safety (AIHS). Except as may be expressly provided by law and subject to the conditions prescribed in the Copyright Act 1968 (Commonwealth of Australia), or as expressly permitted below, no part of the work may in any form or by any means (electronic, mechanical, microcopying, digital scanning, photocopying, recording or otherwise) be reproduced, stored in a retrieval system or transmitted without prior written permission of the AIHS.

You are free to reproduce the material for reasonable personal, or in-house, non-commercial use for the purposes of workplace health and safety as long as you attribute the work using the citation guidelines below and do not charge fees directly or indirectly for use of the material. You must not change any part of the work or remove any part of this copyright notice, licence terms and disclaimer below.

A further licence will be required and may be granted by the AIHS for use of the materials if you wish to:

- reproduce multiple copies of the work or any part of it
- charge others directly or indirectly for access to the materials
- include all or part of the materials in advertising of a product or services or in a product for sale
- modify the materials in any form, or
- publish the materials.

Enquiries regarding the licence or further use of the works are welcome and should be addressed to:

The Manager, OHS Body of Knowledge
Australian Institute of Health and Safety, PO Box 2078, Gladstone Park, Victoria,
Australia, 3043
Manager@ohsbok.org.au

Disclaimer

This material is supplied on the terms and understanding that the Australian Institute of Health and Safety (AIHS) and its respective employees, officers and agents, the editor, or chapter authors and peer reviewers shall not be responsible or liable for any loss, damage, personal injury or death suffered by any person, howsoever caused and whether or not due to negligence, arising from the use of or reliance on any information, data or advice provided or referred to in this publication. Before relying on the material, users should carefully make their own assessment as to its accuracy, currency, completeness and relevance for their purposes, and should obtain any appropriate professional advice relevant to their particular circumstances.

Acknowledgements



The Australian Institute of Health and Safety (AIHS) financially and materially supports the *OHS Body of Knowledge* as a key requirement of the profession.

The *OHS Body of Knowledge* forms the basis of the AIHS OHS capability agenda and informs the other platforms of the agenda: education assurance through accreditation; role clarity, capability assurance through individual certification and continuing professional development.

Thus, the *OHS Body of Knowledge* is strategically important to the AIHS and vital for the profession. (www.aihs.org.au)



The *OHS Body of Knowledge* provides a framework for OHS professional education and continuing professional development. As the body managing accreditation of OHS professional education, the Australian OHS Education Accreditation Board influences, supports and monitors the *OHS Body of Knowledge*, and has a major role in the development and review of individual chapters to ensure that the quality and evidence base reflects current OHS research and leading-edge thinking, and so provides a suitable standard for OHS education and professional development.

www.ohseducationaccreditation.org.au

Bibliography

ISBN 978-0-9808743-2-7

First published in 2019

Author

Professor Sidney Dekker, Director, Safety Science Innovation Lab, Griffith University

This chapter is a companion chapter to OHS BoK 5.1 *Global Concept: Safety* authored by Sidney Dekker in 2012.

Citation of the whole *OHS Body of Knowledge* should be as:

AIHS (Australian Institute of Health and Safety). (2019). *The Core Body of Knowledge for Generalist OHS Professionals*, Ed 2. Tullamarine, VIC: Australian Institute of Health and Safety.

Citation of this chapter should be as:

Dekker, S. (2019). Global Concept: The Science of Safety. In *The Core Body of Knowledge for Generalist OHS Professionals*, Ed 2. Tullamarine, VIC: Australian Institute of Health and Safety.



Global Concept: The Science of Safety

Professor Sidney Dekker PhD

Director, Safety Science Innovation Lab, Griffith University

Email: s.dekker@griffith.edu.au

Sidney gained his PhD from Ohio State University, USA in 1996. In addition to his role at Griffith University in Brisbane, Australia he is a Professor at the Faculty of Aerospace Engineering at Delft University in the Netherlands. As well as his academic roles, Sidney has recently been flying the Boeing 737 as airline pilot and he is a film maker ('Safety Differently,' 2017; 'Just Culture,' 2018) and bestselling author of, most recently: Foundations of Safety Science (2019); The Safety Anarchist (2018); The End of Heaven (2017); Just Culture (2016); Safety Differently (2015); The Field Guide to Understanding 'Human Error' (2014); Second Victim (2013); Drift into Failure (2012); Patient Safety (2011), and co-author of Behind Human Error (2010).

Global Concept: The Science of Safety,

Abstract

While the 'science of safety' is often referred to in professional discussion there has not been any clear exposition of the theories and principles that make up that science. Drawing on the book *Foundations of Safety Science*, this chapter briefly reviews each decade from the 1900s through to the 2010s summarising the predominant safety theories of the time, how they interact and how they are reflected in practice then and now. While the system within which people work is overtly the focus of most of the theories, perspectives from each theory are presented to demonstrate that they actually revert to target the people who work in the system. The chapter concludes that whilst the basis is preliminary and flawed, safety has status as a social science which should inform OHS practice.

Keywords

safety, theory, science, system, people

Contextual reading

Readers should refer to 1 *Preliminaries* for a full list of chapters and authors and a synopsis of the OHS Body of Knowledge. Chapter 2, *Introduction* describes the background and development process while Chapter 3, *The OHS Professional* provides a context by describing the role and professional environment.

Table of contents

1	Introduction	1
2	An episodic view of the development of the science of safety	2
2.1	The early 1900s	2
2.2	The 1910s: Taylor and proceduralisation	2
2.3	The 1920s: Accident-proneness	3
2.4	The 1930s: Heinrich and behaviour-based safety	4
2.5	The 1940s: Human factors and cognitive systems engineering	4
2.6	The 1950s and 60s: System safety	5
2.7	The 1970s and 80s: Man-made disasters, normal accidents and high reliability organisations	5
2.8	The 1980s and 90s: Swiss cheese and safety management systems	6
2.9	The 2000s: Safety culture	7
2.10	The 2010s: Resilience engineering	8
3	People in the system	8
4	Safety as a social science	11
5	A postscript	11
	References	12

1 Introduction

The 'science of safety' is an often used term in discussion among safety professionals but a coherent exposition and understanding of the foundations and interlinking theories that constitute this science has been missing in safety. This has inhibited recognition of safety as a profession and also led to ill-informed discussion, 'knee-jerk' reactions to safety issues and 'flavour of the month' strategies.

The OHS Body of Knowledge has three global concepts: Work, Safety; and Health.¹ Under the global concepts this chapter follows on from OHS BoK 5.1 Global Concept: Safety, written in 2012. In 2012, having been asked to define 'Safety' the author wrote "Safety is a large topic that resists simple definition." (Dekker, 2012, p. 1.) He then explored four questions or viewpoints:

- Is human error a cause or consequence?
- Is compliance with rules a sufficient or limited approach?
- Is safety best conceptualised as absence of negatives or presence of capabilities?
- Is safety best addressed at a component or system level?

This response to the question by posing further questions is, perhaps, an indication of the lack of maturity of professional thinking at the time. Safety has a history. While some 'theories' and practices that emerged through this historical period are now disdainfully rejected by many, other theories are consciously or unconsciously incorporated into our thinking today. But is this rejection/acceptance based on an in-depth understanding of the foundations of the knowledge and a conceptual understanding? This chapter draws on the book *Foundations of Safety Science* (Dekker, 2019) to provide a summary of the evolution of safety theory and how the various theories have influenced our current thinking and practice in safety. It is thus a natural sequel to the dissertation on Safety as a Global Concept written in 2012.

All students of safety and safety professionals should have an in-depth understanding of the theory that underpins their practice. This chapter is a 'global' overview of the evolution of that theory, it is not an in-depth exposition of that theory, that can only be obtained by reading the detail in the *Foundations of Safety Science*. Thus, this chapter does not have detailed referencing.

The core of the chapter is an episodic review of the predominant safety theory of the time with comment on its validity and impact on our thinking today. The threads of the various

¹ See OHS BoK 1.3 Synopsis.

theories are then drawn together to demonstrate that (perhaps unwittingly) almost every approach ends up focusing on the people who work in the 'system'. The chapter concludes by identifying safety as essentially a social science, still in preliminary and flawed, but essential to the practice of safety.

2 An episodic view of the development of the science of safety

Safety science is the interdisciplinary study of accidents and accident prevention. It contains theories inspired by engineering, physical sciences, epidemiology, sociology, psychology, anthropology and more. Most of the theories that guide current safety practices were actually developed during the 20th century.

2.1 The early 1900s

The same principles of scientific experimentation, theorising and logical reasoning—which had increasingly shown their value during the 18th and 19th centuries—could be brought to bear on the problem of safety. This accompanied a broader societal shift. We started believing that the causes of accidents could be scientifically understood, and that there was a moral responsibility to engineer or organise preventative measures. Things didn't just go wrong, accidents didn't just happen because of the will of the divine, or because of fate or random forces in the universe. Things went wrong for a reason. This understanding, and our increasing understanding of how safety was both made and compromised, laid the basis for the emergence of new institutions that could spawn and maintain safety rules and practice. This included regulators, inspectorates and investigative bodies which directly represented the government. Other institutions represented the common interests of employers or workers (standards bodies, professional associations), and still others combined the two, as in the case of government-mandated private insurance schemes. All of these emerged during the twentieth century. The concern with safety was initially politically driven, coming from the mines, factories, railroads, and steamships of the late Industrial Revolution. The interplay between concerned citizens, politicians and scientists, government regulators, and insurers would continue to drive the creation of safety theories and their practical applications into the twentieth and twenty-first centuries.

2.2 The 1910s: Taylor and proceduralisation

Scientific management is an embodiment of the early enthusiasm about being able to control the causes of success and failure, and to carefully manage production and efficiencies. Though the approach stems from the 1910s, it still has implications for how we look at the relationship between work and rules even today. The approach was labelled and exemplified

by Frederick Taylor and the Gilbreths in the first decade (or two) of the twentieth century. They applied a series of scientific tools and methods to determine the most efficient way to perform a specific task. Workers needed to comply with the one best way to do the task in order to maximise efficiency (and safety). Taylorism accelerated and solidified the division between those who plan, manage and supervise the work, and those who execute it. Thinking and problem-solving was heavily concentrated in the former, while the latter merely needed to do what they were instructed. It lies at the basis of the belief that control over workers and operators is possible through better compliance with procedures; through better hierarchical order and the imposition of rules and obedience. Taylor's ideas have left deep traces in our thinking about the safety of work: workers need to be told what to do. They need to be supervised and monitored, and their work has to be specified in great detail. Autonomy and initiative are undesirable. Worker departures from the one best method can be seen as a 'violation,' which requires sanctions or reminders, or the rewriting of the rule (but that shouldn't be done by the worker). They can also be seen as local, adaptive resilience necessary to close the gap between how work is imagined and how it actually gets done. However, adherence to rules can indeed make it impossible to get the work done, lead to blindness to new situations and the squashing of innovation, resentment at the loss of freedom and discretion, and the growth of bureaucratic and supervisory control to impose compliance.

2.3 The 1920s: Accident-proneness

Accident-proneness, a movement that was also quite popular in the first decades of the twentieth century, is one of the earliest attempts to scientifically investigate the 'human element' in safety. It took the emerging sciences of psychology and eugenics, and applied them to explain patterns in industrial data. Accident-proneness built on two patterns that were hard to dispute:

- Human performance is variable, in ways that can be relevant for accidents
- Some people are involved in more accidents than others, in ways that are unlikely to arise purely by chance.

In doing so, it provided an object lesson in the limitations and moral perils of treating human capacities in a reductionist, 'scientific' way. And accident proneness is ultimately difficult to prove. Safety outcomes are never purely due to specific individual attributes. If the same activities take place in different contexts (busy times versus low workload, or night versus day), even the same individual may well have different safety outcomes. It can never be determined with certainty how much of the outcome is attributable to the individual, and how much to context. Today, accident proneness no longer has any status as a 'scientific' theory in safety, even though we might lament certain doctors, drivers or nurses as being more prone to errors (and getting complaints) than others. That said, population-based studies (e.g. male versus female drivers, young versus old drivers) still say something about 'accident proneness,' if anything to support actuarial insurance calculations, but these are not organised around the susceptibility of any particular *individual* within that group to having an accident.

2.4 The 1930s: Heinrich and behaviour-based safety

In the 1930s, Heinrich, an insurance man himself, re-established the idea that many accidents and injuries are preventable. He used a metaphor of a row of dominoes to explain how distant causes lead to injuries. Ancestry and the social environment give rise to character flaws in a worker, such as bad temper, ignorance, or carelessness. Character flaws give rise to unsafe conditions, mechanical hazards, and unsafe acts. These factors in turn lead to accidents, which lead to injuries and fatalities. Like a row of falling dominoes, Heinrich suggested that the sequence could be interrupted by removing the right factor in the sequence. Heinrich's opinion on the best point for intervention shifted throughout his career. Early on, he placed a strong emphasis on improving the physical conditions and physical safeguards of work. Later he placed increasing emphasis on eliminating unsafe acts by workers. He advocated creating an environment where even small undesirable acts are not tolerated. It was mostly through these later, human-error-focused ideas that Heinrich influenced the theory and practice of safety. Behavior-based safety is one of the most visible expressions of it, with us to this day.

Three key ideas of Heinrich's have influenced safety practices (and even some theories) for decades:

- Injuries are the result of linear, single causation
- There is a fixed ratio between accidents (or simply 'occurrences'), minor injuries and major injuries
- Worker unsafe acts are responsible for 88% of industrial accidents.

All three have repeatedly been proven false.

2.5 The 1940s: Human factors and cognitive systems engineering

Despite the three main tenets of Heinrich's work being proven wrong, time and again, he did help us acknowledge how the human became increasingly a recipient of safety trouble—trouble that was created upstream and then handed down to them by their tools, technologies, organisations, working environments or tasks. Human factors took this as its main premise. It was a field that grew from the insights of engineering psychologists in the 1940's. Confronted with the role of the human in increasingly complex technological systems, it represented an important hinge in this thinking about the relationship between humans, systems and safety. Systems and technologies were considered malleable, and they should be adapted to human strengths and limitations. Indeed, individual differences were less important than devising technologies and systems that would resist or tolerate the actions of individuals, independent of their differences. Safety problems had to be addressed

by controlling the technology. The approach of human factors led to a rekindling of our interest in mental and social phenomena. These became important for understanding how best to design and engineer technologies that fit the strengths and limitations of human perception, memory, attention, collaboration, communication and decision making. A few decades later, the field departed from an overly technicalised, individualist, laboratory-task based and mentalist information processing paradigm and took the study of cognition 'into the wild' to understand people's collaborative sense-making in their interaction with actual complex, safety-critical technologies. It led to an entirely new take on human factors, in the field known as cognitive systems engineering.

2.6 The 1950s and 60s: System safety

One offshoot of this, or parallel to it, is system safety. Seeing rapid development in the 1950s and 1960s, it was driven by the commitment that safety should get built into the system from the very beginning. And once a system was in operation, system safety specified the requirements for its effective and safe management. This required system safety to recognise the technical, human and environmental contributors to the creation and erosion of safety, and to map and resolve (to the extent possible) the conflicts and trade-offs between safety and other factors in the design and operation of the system. To do so, systems engineering for safety involves standardised process steps, with many variations in the detail of technique applied at each step. The steps are (semi-) formal modeling of the system under development; analysis of draft system designs; and analysis of the final design to demonstrate safety and to inform post-design safety efforts. From a design perspective, systems can be unsafe through requirements error (designing the wrong system), or implementation error (designing the system wrong). The aim is to prevent foreseeable events and minimise the consequences of unforeseen ones. The increasing complexity of automation and computerisation (particularly when added to legacy systems) can make this very difficult. System safety, through its formal language and techniques, has defined safety as freedom from unwanted events, and protection against unwanted outcomes. As systems have become more complex and anticipating all pathways to failure becomes virtually impossible, the emphasis is shifting to assuring the capacity to handle unforeseen events, rather than assuring the absence of failure modes.

2.7 The 1970s and 80s: Man-made disasters, normal accidents and high reliability organisations

System safety had firmly located safety in the engineering space. There was almost no theorising or practical work on safety outside of it. But a number of high-visibility disasters and a wave of social justice and political movements in the 1960s and 1970s took safety out of the engineering space and its closed, expert-driven language. The size and complexity of many of society's safety-critical systems were becoming apparent to many—and in certain cases alarmingly so. Large disasters with socio-technical systems, and many near-disasters, brought safety and accidents to center-stage (just think of the air disasters of the 1970s and

1980s, as well as nuclear disasters of the era). This greater visibility helped give rise to two decades of productive scholarship, and set the stage for a lot of the conversation about safety, accidents and disasters we are having to this day. Accidents were increasingly understood as social and organisational phenomena, rather than just as engineering problems. Man-made disaster theory was the first to theorise this, closely followed by high-reliability theory and normal accident theory. Disasters and accidents are preceded by sometimes lengthy periods of gradually increasing risk, according to man-made disasters theory. This build-up of risk goes unnoticed or unrecognised. Turner referred to this as the incubation period. During this period, he suggested, latent problems and events accumulate which are culturally taken for granted or go unnoticed because of a collective failure of organisational intelligence. The accumulation of these events can produce a gradual drift towards failure.

This period also gave rise to two (North American) approaches that emerged from the greater social preoccupation with accidents and safety in the 1970s and 1980s. As high-visibility accidents and disasters put the limits of safety engineering and risk management on display, questions arose: Was there a limit to the complexities we could handle? Were there things that we perhaps shouldn't do, or build at all? Normal Accident Theory (NAT) suggested that some accidents are 'normal'—and thus in a sense predictable—because they can be traced to the interactive complexity and coupling of the systems we design, build and operate. Interactive complexity and tight coupling built into the very structure of these systems will generate certain accidents, the theory said, independent of how much risk management we do. Yet there are interactively complex and tightly coupled systems that don't generate accidents, or that haven't yet. So, are there characteristics of 'high reliability' that can somehow be distilled from the things that these systems are doing? This is what High Reliability Organisations (HRO), and High Reliability Theory (HRT) suggested.

A comparison of the NAT and HRT approaches can serve as an introduction to the debate that was triggered by the publication in the early 1990's of *The Limits of Safety* (Sagan, 1993) showed just how close to the edge of total nuclear holocaust the world had actually been on multiple occasions in the decades before (often both triggered and only just prevented by trivial, infinitesimal observations, assessments or actions). Both theoretical schools have had considerable influence on what has happened in safety science since.

2.8 The 1980s and 90s: Swiss cheese and safety management systems

In the late 1980's, the Swiss cheese model became an important icon of the idea that problems experienced at the sharp end (or front line) of an organisation are not created there, but are inherited from imperfect upstream processes and parts. By this time, a strong consensus had already formed that: human performance at the sharp end is shaped by local workplace conditions and distal organisational factors. The Swiss cheese Model is a

defenses-in-depth or barrier model of risk, which suggests (as did Heinrich) that risk should be seen as energy that needs to be contained or channeled or stopped. It's difficult for Swiss cheese to be a true 'systemic' model as it is not capable of explaining or portraying the complex emergence of organisational decisions, the erosion of defenses, drift and a normalisation of deviance. Swiss cheese conceptually aligns with safety management systems. These direct safety efforts and regulation at the administrative end of an organisation, where assurance that safety is under control is sought in management systems, accountabilities, processes and data. The gradual shift to 'back-of-house,' to organisational and administrative assurance of safety, had been long in the making. It is now intuitive that all work is shaped by the engineered context and workplace conditions and upstream organisational factors. If we want to understand or change anything, then that is where we need to look. This trend has also given rise to large safety bureaucracies and cultures of compliance. It has left us with a safety profession that broadly lacks purpose and vision in many industries, and with regulators whose roles have in certain cases been hollowed out and minimised.

2.9 The 2000s: Safety culture

The focus on the sorts of things that can be found and fixed in an organisation before they can create or contribute to an accident has continued, however. Encouraging organisations to build a 'good' safety culture was the logical continuation of this trend, and it spiked in the early 2000s. Safety culture has given organisations an aspiration, getting leaders and others to think about what they want to have rather than what they want to avoid. Researchers and practitioners became concerned with specifying what is necessary inside an organisation and its people to enhance safety. A functionalist approach to safety culture sees and measures it as something that an organisation 'has'. A culture can be taken apart, re-designed and formed. Management can 'work' on parts of that culture (e.g. hazard reporting, procedural compliance). It assumes that values drive people's attitudes and beliefs, which in turn determine their behavior. The interpretivist, or qualitative, approach defines culture as something that an organisation 'does'. It considers culture as a bottom-up, complex, emergent phenomenon; greater than the sum of its parts; resistant to reductionist analysis, measurement, and engineering. For this reason, it cannot be trained and injected into individual minds.

Critiques of 'safety culture' have been targeted at:

- The normative idea that some cultures are 'better' than others
- The implication that cultures are consistent and coherent rather than full of conflict and contradiction
- The avoidance of any mention of power in most models of safety culture
- The methodological individualism that sees culture as the aggregate of measurable individual attitudes, beliefs and values
- The lack of usefulness of the concept in safety regulation and investigations
- The lack of predictive value

- The fact that the ‘container term’ of ‘safety culture’ tries to say so much that it ends up saying very little.

2.10 The 2010s: Resilience engineering

Resilience engineering is the latest school of thinking in safety science. Resilience engineering is about identifying and then enhancing the positive capabilities of people and organisations that allow them to adapt effectively and safely under varying circumstances. Resilience is not about reducing negatives (incidents, errors, violations). Resilience engineering wants to understand and enhance how people themselves build, or engineer, adaptive capacities into their system, so that systems keep functioning under varying circumstances and conditions of imperfect knowledge. How do they create safety—by developing capacities that help them anticipate and absorb pressures, variations and disruptions? Resilience engineering is inspired by a range of fields beyond traditional safety disciplines, such as physical, organisational, psychological and ecological sciences. The organic systems studied in these fields are effective (or not) at adjusting when they recognise a shortfall in their adaptive capacity—which is key to the creation of resilience (or its disruption). Although this is not unique to this theory, resilience engineering, too, appears vulnerable to three analytical traps: a reductionist, a moral and a normative one.

3 People in the system

All approaches in safety science over the past century have seemingly adopted a developmental arc that we would do well to consider, particularly because some practices even today derive from firm and fixed beliefs in the truth of an approach that may be over 40 (Swiss Cheese) or even 80 (Heinrich) years old. What they all have in common, is that from an innovation that typically targeted the system in which people worked, almost every approach seems to end up reverting, one way or another, to the people who work in that system:

- Taylor’s and the Gilbreths’ war on waste and inefficiency analytically deconstructed and then ‘scientifically’ reconstructed the management, organisation and order of work. Prodding individuals to work harder, or smarter, or cheaper, or better, or safer was of no use. The tasks they were told to do had to be strictly systematised first, and the tools with which they were told to work needed to be methodically engineered.
- Epidemiological data started pointing to a differential probability of suffering harm and accidents, which eventually led to the accident-prone thesis. The original impetus was systematic and scientific, however, and its interventions were targeted at the level of the system in which people worked, matching skills with demands.
- Heinrich was one of the first to systematically investigate ways to stop hazard trajectories from causing accidents and injury. On practical grounds, he advocated

selecting remedies as early as possible—upstream, in the selection of work and workers, in the design of work and workplaces. He placed a strong emphasis on improving the physical conditions and physical safeguards of work.

- Human factors was born out of the realization that the human was the recipient of error-prone and error-intolerant system designs. Operational people inherited safety problems (from their technologies, tools, organisations, working environments and conditions) rather than creating safety problems. Solutions were targeted at the technology, the system, which were systematically adapted to human strengths and limitations.
- System safety promoted the notion that safety needs get built into the system from the very beginning. And once a system is in operation, system safety needed to specify the requirements for effective and safe management of the system. It shouldn't be left to frontline heroics to make things work in practice, or to recover from built-in error traps. These could, and should, be designed and managed out—upstream.
- Man-made disaster theory understood accidents and disasters squarely as administrative or organisational phenomena. The intentions and actions that incubated accidents and disasters were bred—by very normal, everyday processes and bureaucratic workings—at the organisational blunt end. Prevention efforts should be targeted there—upstream.
- Normal Accident Theory proposed how structural features of a system—its interactive complexity and tight coupling—not only create fundamental paradoxes in any attempts to safely manage it, but set that system up for a predictable kind of breakdown: the normal accident. The target for safety improvements is not in the people managing or operating it, but in the system itself and in the political arena that allows it to operate at all.
- Reason showed that the further people were removed from the front-line activities of their system, the greater the potential danger they posed to it. Merely responding to front-line errors or violations would not lead to improvements. Attempts to discover and neutralise 'resident pathogens' in the system, the organisation, was going to have a greater beneficial effect on system safety than chasing localised acts by those on the sharp end.
- Safety culture gave organizations an aspiration: it got leaders and others to think about what they wanted to *have* rather than what they needed to *avoid*. The target was more than just upstream, it encompassed the entire organisation and its culture. Researchers and practitioners became concerned with specifying what is necessary inside an organisation to enhance safety.
- Resilience engineering represents a valiant and honest attempt to specify the possibilities for creating safety in complex, dynamic systems. In these systems, novel phenomena emerge, and much more variation occurs than could ever be procedurally specified, or designed or trained for. Engineering resilience into them is a matter, in part, of recognising and matching the requisite variety of the system.

The innovations that have come from these approaches gradually congealed into their respective sets of technocratic projects, of programmatic operations and tools. The approaches all successively became professionalised, spawning time-and-motion researchers, psycho-technologists, behavioral safety consultants, human factors engineers, system safety experts, safety professionals and practitioners trained in methods such as STAMP and FRAM and ICAM.

Almost invariably, in this transformation of the approach—that started with the system, the technology, the organisation, upstream—it lands back on the individual people in that system. In so many words, it appeals to these people to try yet a little harder. The original impetus targets the system; the eventual appeal is directed at the individual:

- Taylor's call to become a 'high-class man' could be answered only by an individual submitting completely to the emptiness of the tasks designed and 'scientifically managed' by others. This kind of work is now known to have been so mind-numbing or soul-destroying that Henry Ford had to 'bribe' his workers with a \$5-a-day wage.
- Accident-prone theory descended into the morally dubious (and scientifically untenable) separation of workers who were 'fit' from those who weren't, justifying why we could legitimately give up hope on some people; and aligning itself with theories and practices that quantified who was an idiot, a moron, an imbecile.
- Heinrich placed increasing emphasis on eliminating unsafe acts by workers. He advocated creating an environment where even small undesirable acts were not tolerated. It spawned behavior-based safety. This squarely targeted workers and at times ignored the work or the environment of that work. It can deteriorate into retributive rituals in response to not meeting low-injury targets and arbitrary firings or exclusions on account of 'safety violations.'
- From a focus on solving safety problems by targeting the technology and the system, human factors has been enjoined to revisit the individual. It has been co-opted into methods and theories that squarely target human shortcomings, such as line management deficiencies, supervisory shortcomings, fallible decisions, unsafe acts and violations. In some accident reports, even the dead can now be blamed for their own complacency and 'loss of situation awareness.'
- When pushed to explain how disaster incubation happened, man-made disasters had little else to go on than collective human deficiencies—erroneous assumptions, not noticing information or misunderstanding signals, communication failures, not heeding warnings, and a human reluctance to imagine the worst.
- The structural analysis of Normal Accident Theory was counterbalanced with high-reliability research that showed the extent to which people actually go (and go successfully) to manage even seemingly unmanageable complexity. But this too, hid a moral judgment. The recent focus on 'mindfulness' allows this same approach to say that failure equals a lapse in detection, that someone somewhere didn't anticipate or catch what went wrong, and that they should have paid attention earlier.
- From an enthusiastic embrace of systems thinking, Swiss cheese became entwined in a new effort to focus on the individual. In our search for distal causes, we had thrown the net too widely; we let the pendulum swing too far toward the system to be

morally justifiable. The human could occasionally be a hero, but was best treated as a hazard, a system component whose unsafe acts are implicated in the majority of catastrophic breakdowns.

- Safety culture research, when operationalised, stopped being about culture and upstream organisational factors fairly swiftly. It instead became reduced to the attitudes and behaviors of individuals in the organisation. Safety culture could be improved by targeting those in ways that are similar to what behavior-based safety might propose (hearts-and-minds campaigns, posters, incentive programs—all focusing on fixing the worker, not the work or the environment).
- From an explicit commitment to understanding some of the complexity and variability of the systems in which people work, and literally helping to engineer resilience into these systems, resilience engineering has been turned around and used to blame people for not having enough resilience. Its approach has been deployed to force individuals to accept and adapt to dangers that brew and grow beyond their control.

4 Safety as a social science

At the heart of this pattern is a dialectic—system or person, upstream or downstream, organisation or individual? Of course, we should look at the system in which we and other people work, and improve it to the best of our ability. At the same time, we know that all safety-critical work is ultimately channeled through relationships between human beings, or through direct contact of some people with the risky technology. A system creates all kinds of opportunities for action. And it also constrains people in many ways. Beyond these opportunities and constraints, we could argue that there remains a discretionary space, a space that can be filled only by an individual care-giving or technology-operating human. This is a final space in which a system really does leave people freedom of choice (to launch or not, to go to open surgery or not, to fire or not, to continue an approach or not). It is a space filled with ambiguity, uncertainty and moral choices. This is also why people matter: they will forever be (somewhere) behind the design and operation of these systems, if not actually actively operating them. That is why safety science will forever be, in an important sense, a social science. The future for safety science may well lie in our ability to break out of this dialectic and see people *in* systems, rather than people versus systems.

5 A postscript

There are organisations that have developed safety into something akin to a ‘mystery religion,’ replete with its sacred texts, saintly thought leaders, holy relics, sacred places (think of ten golden rules, high-viz vests and altar-like safety notice boards), rituals (the prayer-like ‘safety share’) and idolatry: the giving of ultimate meaning and allegiance to something that should be serving only an intermediate purpose (like Zero Harm, which is at best a moral mobiliser, an inspiration or aspiration, but is idolised as a promised land that

will deliver mankind from suffering). But safety is not a mystery religion. It is as pedestrian and flawed and preliminary and human as any science, as any operational practice. But there is a science to safety and it is only by recognising, examining and objectively testing the science that we will advance the science of safety and so the safety of workers and others.

References

Dekker, S. (2019). Global Concept: Safety. In *The Core Body of Knowledge for Generalist OHS Professionals*, Ed 2. Tullamarine, VIC: Australian Institute of Health and Safety.

Dekker, S, (2019). *Foundations of Safety Science: A Century of Understanding Accidents and Disasters*. Boca Raton, FL: CRC Press.

Sagan, S. (1993). *The Limits of Safety: Organizations, Accidents and Nuclear Weapons*. Princeton, NJ: Princeton University Press.