Essay-1: From Resilience Engineering (RE) to Incremental Safety (IS):

Precis: Resilience Engineering is obviously the intellectual progenitor of incremental safety. This essay recounts how resilience engineering came about, and also tries to explain why resilience engineering did fullfil the expectations by which it was met, specifically to become the long awaited replacement of the Safety Legacy (SL). RE was the reason for proposing Safety-I and Safety-II, and since Safety-II now is morphing into incremental safety, RE is clearly also part of the foundation for incremental safety.

Resilience Engineering (RE) became part of safety discussions, following a first symposium on the topic organised in Söderköping, Sweden in November of 2004. The most tangible result of that symposium is a slowly growing number of books about resilience engineering, (Hollnagel, Woods & Leveson, 2006), (Hollnagel, Nemeth, & Dekker, 2008), (Nemeth, Hollnagel, & Dekker, 2009), (Hollnagel, E., et al., 2011), and (Nemeth, & Hollnagel, 2022), as well as the creation of the Resilience Engineering Association (https://www.resilience-engineering-association.org/). The primary purpose of the symposium was to formulate the basis for Resilience Engineering. Needless to say, none of the invited participants in 2004 including the symposium organisers had a clear idea about what resilience was let alone what Resilience Engineering (RE) should become. The participants shared the need to find something that could disrupt the existing safety legacy impasse, best illustrated by the fact that all safety statistics, which all are measures of the absence of safety, eventually reach a point of diminishing return where further efforts to reduce the number of accidents have no measurable effect but seem to reach an asymptote, Figure 1, below is a characteristic example of that.

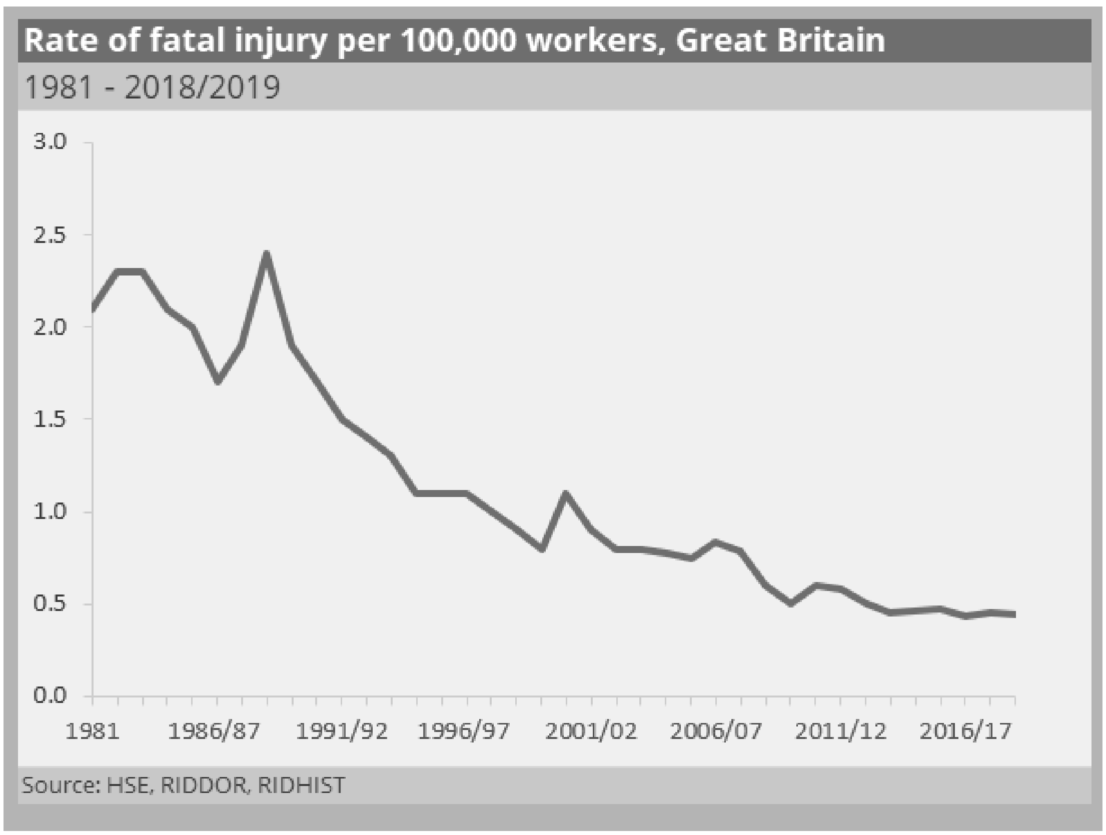

Figure 1.1: Rate of fatal injury per 100.000 workers (Great Britain) (Health and safety statistics, House of Commons Library

Throughout these essays I intend to avoid the use of the term accident, and will instead write Unexpected and Unacceptable Outcomes (UUOs) and in most cases just use the abbreviation UUO. An UUO precisely describes what an accident is: Unexpected when it happens, with unacceptable consequences for whoever is present. This rendering actually fits well with the origin of the term. The word accident comes from the Latin *accidere*—a combination of *ad-* ("to" or "towards") and *cadere* ("to fall"). It originally meant "an occurrence or event that comes by chance," and literally translated to "falling upon." And this perfectly matches the meaning of an UUO. Few people willingly stumble and fall, circus clowns excepted. Only people with ill intentions, either a terrorist or an arsonist may want outcomes that are harmful to people, but in these cases the harmful results are EAOs for the perpetrator, and UUOs for their victims.

The use of the UUO and EAO terms conveniently overcomes the problem that the language used by the safety legacy does not have an antonym for accident, which we actually need to refer to what we want. The vocabulary of the safety legacy is rich as far as negative stance is concerned with many terms for the UUOs, and especially their putative causes, and correspondingly poor for EAOs. The opposite of an UUO is an Expected and Acceptable Outcome (EAO), and to ensure EAOs agree with most production and quality departments. The terms UUO and EAO can therefore in their own small way help to dismantle the conventional organisational silos, and thereby become a step toward *synesis* (Hollnagel,2021). The remit of a safety department is to reduce the number of UUOs, which, of course logically can happen indirectly by increasing the number of EAOs (an outcome can neither be expected and unexpected at the same time, nor acceptable and unacceptable. To increase the number of AEOs, however, corresponds to the remit of production and quality departments: An investment to further safety will therefore also be an investment that furthers productivity. And who could object to that? (It may be objected thast this is mostly a semantic trick, yet it is difficult to recognise this truly synergetic relation unless a useful antonym for accident is available.

Due to the safety legacy impasse, anything different was welcome even though in 2024 few knew much about the history and etymology of resilience. That changed as soon as resilience became popular. According to (McAslan, 2010), the modern understanding of resilience, can be traced back to the Canadian ecologist Crawford Stanley Holling who applied it to characterise ecological systems (Holling, 1973). In relation to safety, the term was first used by (Woods, 2000). "Resilience" again has its origin in a Latin word *Resilio*. (*re* means "back", and *silio* means "to leap") (Hosseini, Barker & Ramires-Marques, 2016). The original meaning described the deformation and rebound state of objects. Physicists first used it to express a spring's characteristics and elaborate on a material's stability against external shocks, whether it was able to leap back. Around 1850, the Irish geophysicist, civil engineer and artillery designer Robert Mallet used resilience as a way to compare the strength of materials used in the construction of guns for the Royal Navy's war ships. Mallet used the concept to explain why "in bronze guns the expansion is so great and the resilience, or power of elastic recovery, so small that in extreme cases the gun becomes permanently lengthened. Later, the label resilience attracted the attention of scholars in many other fields. Predictably leading to more or less sensible combinations such as ecological resilience, infrastructure resilience, city resilience, organisation resilience, economy resilience, psychological resilience, social resilience, and several others, including cyberresilience. RE was bid welcome as something that possibly might get the safety legacy out of the stalemate it had reached. RE was, however, from the very start stuck knee-deep in the safety legacy conceptual mud: (1) in terms of the dichotomy between what goes well and what fails (the success / failure, here called expected and acceptable outcomes (AEO) and unexpected and unacceptable outcomes (UUO) respectively) and (2) in terms of the **negative stance** represented by the efforts of distancing and avoidance in both learning and safety management. The safety legacy focus of learning (Essay-5) is on what **not to do** on what to *prevent*, *weaken,* and *eliminate*, and learning was traditionally based on the occurrence of unexpected and unacceptable outcomes, (Kletz 2001), although some (Anand, 2024) argue we are not learning effectively. It obviously goes against common sense to propose that we can learn about safety from conditions where safety is absent. It clearly makes more sense to learn about safety from conditions where it is present, cf., the first paradox of Reason (Essay-2). For safety management the aim of the safety legacy is to ensure that there are as few UUOs as possible, (corresponding to Zero Accident Vision (ZAV)), and the safety legacy embodies a juxtaposition between work that goes well and work that fails, with almost all emphasis on the latter. Resilience rightly aimed to go beyond the ability to recover from disruption and diversity and beyond notions of graceful degradation or graceful extensibility.

The safety Legacy

The safety legacy comprises the widely and uncritically accepted assumptions about how UUOs happen and how events that lead to such outcomes develop. The safety legacy is in short the underlying, unspoken assumptions that determine how we perceive and interpret the occurrence of UUOs, and therefore also how we respond to them -- what we do about them in both the short and the long term. A legacy is either something that intentionally is left or bequeathed to someone else, or the unintentional long lasting effect of something that has happened. An intentional legacy is usually positive and welcomed by those who receive it, while the opposite is the case for an unintentional legacy. The safety legacy is definitely of the latter type and comprises the conglomerate of assumptions and pet hypotheses that constitute the language and terminology safety professionals, experts, consultants, and managers rely on when they try to determine socially acceptable causes for UUOs, and when they communicate their conclusions to colleagues, and / or the general public as well as politicians and other official decision makers. The safety legacy was never formally announced or put forward as a legacy, in the way that (Heinrich, Petersen & Roos, 1980), proposed their ten axioms of safety, or that James Reason much later described the four paradoxes of safety.The safety legacy is in turn derived from the **Heinrich dogma** (Hollnagel, 2025), which essentially is the first of ten axioms of safety later proposed by (Heinrich, Petersen & Roos, 1980, p. 21), which goes like this:

> "The occurrence of an injury invariably results from a completed sequence of factors—the last one of these being the accident itself. The accident in turn is invariably caused or permitted directly by the unsafe act of a person and/or a mechanical or physical hazard."

Some companies use safety legacy in their promotions, and define it, for instance, as ”a priority on people and values, and on strong safety cultures that extend beyond compliance. Our mission is to ensure everyone goes home safe every day.” (But this curiously disregards possible accidents after people begin their journey home, and before they arrive at home!) This is obviously not the responsibility of the company they work for, but nevertheless illustrates how the safety legacy tempts us to decompose issues into constituent parts, and solve them in isolation). For the loved ones it is equally tragic if someone dies at work or in a traffic accident on the way home. It is common knowledge that airline pilots are subject to greater risk, when they jetlagged drive home from the airport than during a flight,e.g., between the US and Europe. The main issue with the legacy is that it embodies, and therefore gradually, and surreptitiously introduces a habit of simple linear cause-effect reasoning, perhaps best illustrated by the Domino model.The legacy also commonly endorses very generic solutions, such as safety culture (Essay-4). other popular models of the safety legacy, such as the Swiss Cheese Model (SCM) (Reason, Hollnagel, & Pariès, 2006) and the Bowtie model (de Ruijter, & Guldenmund, 2016) also rely on reasoning by linear progression, the SCM, by how the active failure or harmful influence slowly penetrates the layers of defense (the slices of cheese), and the Bowtie, because the Bowtie is actually just a combination of two components the fault tree and the event tree that each is linear by nature. The proper technical name for the combination of a fault tree and an event tree is anyway Cause-Consequence Control or CCC (Taylor, 1974), with no reference to the visual outline of a bowtie. (Figure 1).

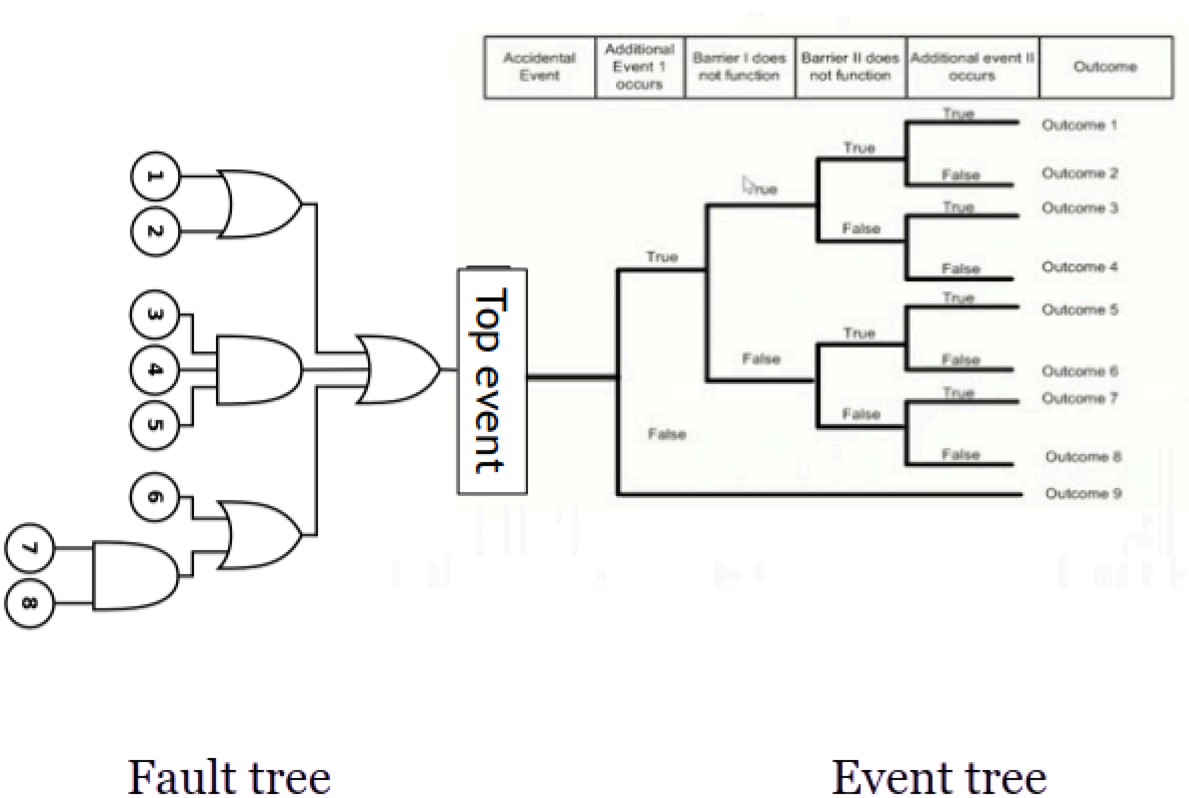

Figure 1: A simple cause-consequence tree (Bowtie prototype)

RE never offered a simple ready-made model of its own that could overcome the limitation imposed by linear causal reasoning. There exists a method that offers an alternative way, the Functional Resonance Analysis Method (FRAM) (Hollnagel, 2012) and (Patriarca, 2026). Although the FRAM was developed explicitly to overcome the limitations of linear causal thinking, it was never intended to become a method for RE, that it is seen so by many anyway is mostly a temporal coincidence.

Resilience is not a unitary system quality, even though it often is treated as such. Resilience can therefore neither be engineered nor measured as such. A system cannot be, and cannot have resilience, but a system can perform in a way that is resilient. Resilient performance represents an ongoing condition where problems are momentarily under control due to compensating changes. Being able to do so is necessary in order to function reliably in environments where unexpected and unpredictable changes are possible and where their consequences can propagate rapidly. To perform resiliently, a system must have the potential to respond to the actual, the potential to monitor what may become critical, the potential to learn from the factual in particular from work that went well, and the potential to anticipate the potential, future changes that may affect any of the potentials.

Two key threads that were present from the start continue to inform Resilience Engineering (RE).

- People create safety and performance variability is an indispensable asset rather than a liability. Where the safety legacy traits people as a liability, as error-prone and unreliable. RE acknowledges that people are an indispensable asset, and that variability frowned upon by the safety legacy, is agility rather than variability. This had already been expressed by the third of the four safety paradoxes described by professor James Reason, listed below. Human performance variability ought to be called human **performance agility**, since it effectively is the putty that fills out the cracks due to the difference between Work-as-Imagined(WAI) andWork-as-Done (WAD).
- Adaptive capacities need to be developed and supported because they are critical in order to overcome risks from brittleness given increasing complexities (now better expressed as complexity effects that lead to brittle systems; adaptive capacities are essential for systems to remain competent and be able to develop further) (Lay & Branlat, 2022).

The common use of the term resilience, within safety and elsewhere, retained a negative connotation because it focused on how a system handled diversity, stresses, and disruptions. The negative connotations were understandable considering that the origin of resilience, or the first modern uses of the term, was in physics and material science. Since a physical material is passive, and can only respond, resilience must be seen in relation to, or as a reaction to, the potentially harmful or destructive consequences of unexpected events. Even when the idea of resilience was picked up by ecology, the focus on adversity remained. Although an ecological system is dynamic where a physical system is static, it is still not conscient. It remains reactive, perhaps with an element of random responses thrown in, and only acts in response to something that happens, to something that in one way or another is imposed by external forces or agents.

When resilience came into the field of industrial safety (in a wide sense), as in resilience engineering, the negative and reactive connotations remained. That is not hard to understand either, since safety traditionally has been preoccupied with avoiding adversity, risks, and harm. This is nicely captured by the first of the four safety paradoxes described by James Reason:

"Safety is defined and measured more by its absence than its presence." (Reason, 2000, p. 1).

The focus on the negative, on sustaining an organisation despite adversity, is very common and resilience is therefore seen as the ability of an organisation to react to and recover from disruptions with minimal effect on its dynamic stability.

Resilience engineering is, however, not only about dynamic systems but about socio-technical systems, about organisations as deliberate or intentional configurations of people, materials, and activities (and information) that have been conceived to achieve a given objective or objectives. It is therefore insufficient to think of resilience *vis-a-vis* the negative. In order for an organisation to exist it must not only respond **when** something happens but also act **before** something happens. not only in the face of dangers, trying to protect itself, but also in the face of opportunities that allow it to survive and grow – in every possible way. Recognising and responding to opportunities is indispensable for everyday activities throughout an organisation for individuals, for social groups (collectives), for management, and for the organisation itself. There is a similarity to ecological systems, which also grow when the opportunity arises. The important difference is that ecological systems neither look for opportunities nor anticipate opportunities. Nor do they try to create or bring about opportunities. But organisations do. Organisations are – and have to be both – strategic and tactical.

Accentuate the Positive

As the above brief history shows, the thinking about resilience has typically referred to a dichotomy: on the one hand materials, systems, or situations where resilience was absent and where adverse outcomes therefore might result, and on the other hand materials, systems or situations where resilient was present and where adverse outcomes could be avoided. This was also the case in the early 2000s when resilience engineering was proposed as an alternative (or rather as a complement) to the conventional view of safety. The first book provided the following definition.

> "The essence of resilience is therefore the intrinsic ability of an organisation (system) to maintain or regain a dynamically stable state, which allows it to continue operations after a major mishap and/or in the presence of a continuous stress." (Hollnagel, 2006).

The definition reflected the historical context by juxtaposing two states - one of stable functioning and one where the system had broken down. But this also limited the definition to consider situations of threat, risk or stress.

Discussion about resilience versus robustness or resilience versus brittleness nevertheless soon made clear that resilience is not only about avoiding failures and breakdowns and that resilience is not just the converse of a lack of safety. In a later book, book, the definition became:

> "The intrinsic ability of a system to adjust its functioning prior to, during, or following changes and disturbances, so that it can sustain required operations under both expected and unexpected conditions."

(Woods & Branlat, 2011, p???. Hollnagel et al., 2011, p. 1.).

In this definition the reference to risks and threats had been replaced by a reference to 'expected and unexpected conditions'. The focus had also changed from 'maintaining or regaining a dynamically stable state' to the ability to 'sustain required operations'. The logical continuation of these developments leads to the following definition:

> Resilience is an expression of how people, alone or together, cope with everyday situations – large and small – by adjusting their performance to the conditions. An organisation's performance is resilient if it can function as required under expected and unexpected conditions alike (changes / disturbances / opportunities). (Hollnagel, 2017, p. 20).

The changes in the definitions broadened the scope of resilient performance. It was not just about being able to recover from threats and stresses, but rather about being able to perform as needed under a variety of conditions – and to respond appropriately to both disturbances and opportunities. The focus of resilience engineering is thus ***resilient* performance**, rather **resilience** as a property (or quality), hence the adverb *resilient* rather than the noun *resilience*.

The emphasis on opportunities is important for the change from decremental, and protective safety (Safety-I) to incremental, and productive safety (Safety-II) – and ultimately for the dissociation of resilience from safety, thereby leaving the sterile discussions and stereotypes of the past behind. Resilience is about how organisations perform safely (Hollnagel, 2025), not just about how they remain safe. No organisation or company actually manages safety, they manage how well they do what they are supposed to do since that is what their customers pay for. An airline may well have a Safety Management System (SMS), but an airline manages the flights it offers, so they are comfortable and reliable, since that is what passengers are willing to pay for, they obviously also try to avoid regular type accidents(Weatrum, 2006), because they may be costly in terms of material damage and loss of reputation. It is good business to avoid unexpected and unacceptable outcomes. Business class and economy always differ in terms of comfort, but rarely in terms of safety.

An organisation that is unable to make use of opportunities is in no better position than an organisation that cannot respond to threats and disturbances – at least not in the long run.

The purpose of resilience engineering is to ensure that an organisation can perform effectively in unpredictable everyday conditions, in other words do everyday work successfully. It is also about ensuring that an organisation is able to cope – the unexpected situations – both when they have the potential to disrupt functioning or performance (**threats and risks**) and when they have the potential to improve or augment everyday performance (**opportunities**). Indeed, resilience is not only found during non-routine or critical incidents but also, and perhaps more so, when incidents are not happening, i.e., during routine functioning with acceptable outcomes, that (Weick, 1987, p. 118) so elegantly called "dynamic non-events" Ecology and business agree that the resilience of an organisation is an expression of its ability to sustain its own existence – to survive and to thrive. The survival or continued existence is in itself a proof of resilience, but not synonymous with it. Resilience succeeded in becoming a popular new term of the safety vocabulary, but resilience engineering never replaced the safety legacy. One reason why was because it retained the two major weaknesses of the safety legacy mentioned above, (1) the dichotomy between what goes well and what fails (success / failure, here called EAOs and UUOs, respectively) and the negative stance. A second reason was that RE did not offer any accident models or investigation methods that effectively could rival the domino model and RCA. These are the challenges that incremental safety must overcome, with the benefit of the experiences from RE.

Resilience versus brittleness

Resilience is the intrinsic ability of an organisation to adjust its functioning prior to, during, or following changes and disturbances, so that it can sustain required operations under both expected and unexpected conditions (Hollnagel, 2011). Resilient systems are agile in the face of change and have buffers (margins of maneuver) to respond to unforeseen demands, thus creating the conditions needed to avoid or minimise consequences of adverse events; they understand and manage their complexity. Brittle systems, on the other hand, fail to notice warnings and to adjust their behavior in time to prevent collapse; they overlook and fall victims of tight couplings (Perrow, 1984). They may have a system designed around "standard" (low variability) maintenance even though variable scope is the norm (standard conditions are improbable).

Brittleness, whether of a social system or a physical artefact is obviously not an attractive quality. It is an indisputable advantage if both social systems and physical artefacts are resilient (Figure 2). But that does not mean that resilience can be defined as the opposite of brittleness. It goes against any kind of scientific reasoning to define something as the opposite of something else. A definition shouldexplain what something is, rather than what **it is not.**

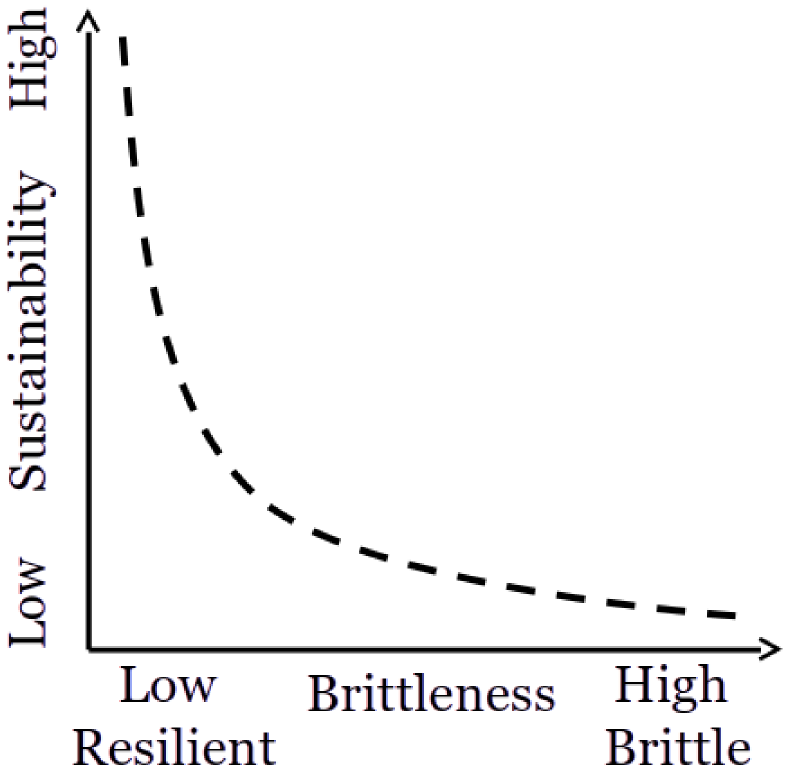

Figure 2: The relation between sustainability and brittleness

Brittleness and resilience are not just two sides of the same coin as (Cook, 2012) made clear. They are rather two different and incompatible qualities of a system or a system or an artefact. Just as where there will be risks and threats, there will be chances and opportunities; where an organisation or system (or process) is brittle, the possibility to increase resilience exists. Brittleness and resilience are system properties not outcomes (Cook, 2012). You can have an acceptable outcome from a brittle system or an unacceptable outcome from a resilient system although the probability of having a desired outcome is larger for resilient systems. The "Resilience Engineering…agenda is to control or manage a system's adaptive capacities based on empirical evidence." (Lay & Branlat, 2014, p. 142). "To achieve resilient control…a system must have capacity to reflect on how well it is adapted, what it is adapted to, and what is changing in its environment."(Ibid, p. 144). Managers need knowledge of how the system is resilient and brittle to make decisions on how to invest resources to increase resilience (Woods, 2006; Hollnagel, 2009). "Resilience / brittleness of a system captures how well it can adapt to handle events that challenge the boundary conditions for its operation." (Woods and Branlat, 2011). Brittle systems and brittle equipment may malfunction or break. These are obviously undesired outcomes. Resilient systems and resilient equipment are less likely to malfunction or break. But this likelihood is not a measure of resilience as such.

Why did RE not succeed?

Looking back at the first symposium on RE in 2004 and what followed, it is clear that RE did not manage to replace the safety legacy or even to bring about any significant changes in the practice of safety as safety investigations and safety management. It is worth spending a little time to discover why it failed. Several of the reasons have already been mentioned.

One reason is that RE from the very start stuck knee-deep in the safety legacy's conceptual mud and that it regrettably and unintentionally continued the negative stance, as seen by the emphasis on brittleness at least during the first several years. Another, and probably more important reason is the RE never was able to propose or develop ready-made methods that practically could introduce RE principles and practices to organisations, perhaps because there never were any clearly formulated principles or practices. From a historical perspective changes to models and methods have mostly been a consequence of changes in thinking about safety, for instance the three ages described by (Hale & Hovden, 1998). These changes in thinking in turn became necessary, because hitherto unknown types of UUOs occurred. RE came about because of a gradually growing dissatisfaction with the safety legacy, and that was not by itself sufficient to bring about a new age in safety thinking. Neither were ever there any ready-made resilience models or methods to rival the oversimplified accident models and investigation methods that the safety legacy provided. There was nothing practical or operational that could be used directly.

The first paradox is that "Safety is defined and measured more by its absence than its presence." (Reason, 2000, p. 1.) is amply illustrated by the common definitions of safety. The American Society Of Safety Engineers, for instance, defines safety as “the freedom from unacceptable risk.” (ASSE, 2011 p.13) The focus of practical efforts therefore is how best to reduce the number of unexpected and unacceptable outcomes to an acceptable level (ALARP), ideally zero, and the emphasis is on how to manage these, as demonstrated by the ubiquitous safety management Systems (SMS). Similarly, the International Civil Aviation organisation (ICAO, 2013) defines safety as ”the state in which the possibility of harm to persons or of property damage is reduced to, and maintained at or below, an acceptable level through a continuing process of hazard identification and safety risk management. (ICAO, 2013p. 1-2.).

References

Busch, C. (1019). Heinrich’s Local Rationality: Shouldn’t ‘New View’ Thinkers Ask Why Things Made Sense To Him? Lund, Sweden: Thesis/Project work submitted in partial fulfilment of the requirements for the MSc in Human Factors and System Safety.

Björnberg, K. E., et al. (2019). *The vision zero handbook*. Cham, Switzerland: Springer International Publishing.

Cook, R. I., (2012) *Why resilience matters?* Presentation at University of BC School of Population and Public Health Learning Lab, May 7

Cook, R. I., & Nemeth, C. (2006). Taking Things in One’s Stride: Cognitive Features of Two Resilient Performances. In: E. Hollnagel, D. D. Woods, & N. G. Leveson (Eds.), *Resilience Engineering: Concepts and Precepts* (pp. 205–221). Adelshot, UK: Ashgate.

De Meyer, A., Loch, C. H., Pich, M. T. (2002). Managing Project Uncertainty: From Variation to Chaos, *MIT Sloan Management Review*, Winter Vol., 60 - 67.

de Ruijter, A. & Guldenmund, F. (2016). The bowtie method: A review. *Safety science*, *88*, 211-218.

Holling, C. S. (1973) Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics*, *4*, 1-23.

Hollnagel, E. (2009). The ETTO Principle: Efficiency-Thoroughness Trade-Off-Why Things That Go Right Sometimes Go Wrong. Farnham, UK: Ashgate.

Hollnagel, E. (2012). *The Functional Resonance Analysis Method: Modelling complex socio-technical systems*. Farnham, UK: Ashgate.

Hollnagel, E. (2017). *Safety-II in practice*. Abingdon, Oxon, UK: Routledge.

Hollnagel, E. (2011) Epilogue in: E. Hollnagel, et al., (Eds.) *Resilience engineering perspectives, Vol 3.: Resilience engineering in practice*. Farnham, UK: Ashgate.

Hollnagel, E., Woods, D. D. & Leveson, N. G. (Eds.), (2006). Resilience engineering: Concepts and precepts. Aldershot, UK: Ashgate.

Hollnagel, E., Nemeth, C. P. & Dekker, S. W. A. (Eds.), (2008). Resilience engineering perspectives, Volume 1: Remaining sensitive to the possibility of failure. Aldershot, UK: Ashgate.

Hollnagel, E., et al., (Eds.), (2011). *Resilience engineering perspectives, Vol 3.: Resilience engineering in practice*. Farnham, UK: Ashgate.

Hosseini S, Barker K, Ramires-Marques, J. E. (2016).A review of definitions and measures of system resilience. *Reliability engineering & system safety, 145*: 47-61

Hong, L. & Page, S. (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. PNAS, Vol. 101, 46 (16385-16389)

Lay, E. (2011). Practices for Noticing and Dealing with the Critical. A Case Study from Maintenance of Power Plants. In E. Hollnagel, J. Pariès, D. D. Woods, & J. Wreathall (Eds.), Resilience Engineering in Practice (pp. 127–144). Farnham, UK: Ashgate.

Lay, E. & Branlat, M. (2014). Noticing Brittleness, Designing for Resilience. In:C.P. Nemeth & E. Hollnagel (Eds.) *Resilience engineering in practice, Volume 2. Becoming resilient*. Farnham, UK: Ashgate.

Lay, E & Branlat, M.(2022). Noticing Brittleness, Designing for Resilience. In: C. P. Nemeth, & E Hollnagel (Eds.) *Advancing resilient performance*. Springer International Publishing (142, 152).

McAslan, A. (2010). The concept of resilience: Understanding its origins, meaning and utility. Adelaide: *Torrens Resilience Institute,* 1.

Nemeth, C. P., Hollnagel, E. & Dekker, S. W. A. (Eds.), (2009). *Resilience engineering perspectives,Volume 2: Preparation and restoration*. Farnham, UK: Ashgate.

Nemeth, C. P., & Hollnagel, E. (Eds.). (2022). *Advancing resilient performance*. Berlin, Heidelberg: Springer International Publishing.

Nielsen, D. S. (1970). *The Cause Consequence diagram method as a basis for quantitative accident analysis, (Report Risø-M-1374)*. Roskilde, Denmark: Risø National Laboratory

Patriarca, R. (Ed.)(2026). *Navigating the FRAM Mastering the Functional Resonance Analysis Method for Modelling Complex Socio-Technical Systems.Abingdon*. Oxon, UK: Routledge.

Perrow, C. (1984). *Normal Accidents: living with high-risk technologies*. Princeton, NJ: Princeton University Press.

Taleb, N. N. (2010). *The Black Swan*. 2nd edition. Random House.

Westrum, R. (1993). *Thinking by groups, organisations, and networks: A sociologist's view of the social psychology of science and technology.* In W. Shadish, & S. Fuller (Eds.), The social psychology of science (pp. 329-32). New York: Guilford.

Platt, M. L. & Huettel, S. A. (2008). Risky business: the neuroeconomics of decision making under uncertainty. *Nature Neuroscience*, *11*, 398 – 403.

Reason, J. T., Hollnagel, E., & Pariès, J. (2006). Revisiting the Swiss cheese model of accidents. *Journal of Clinical Engineering, 27*(4), 110-116.

Taylor, J. R. (1974). *Sequential Effects in Failure Mode Analysis(R-12-74)*. Roskilde, Denmark: Risø National LanoratoryLaboratory, Electronic Department (R-12-74)

Woods, D. D. (2000). *Designing for resilience in the face of change and surprise: Creating safety under pressure*. Plenary Talk, Design for Safety Workshop, NASA Ames Research Center, October 10.

Weick, K. E. (1987). Organizational culture as a source of high reliability. *California Management Review, 29*(2), 112-128.

Weick, K. E., & Sutcliffe, K. M. (2001). *Managing the Unexpected: Assuring High Performance in an Age of Complexity* (1st ed.). Jossey-Bass.

Westrum, R. (2006). A Typology of Resilience Situations. In Hollnagel, E., Woods, D. D. and Leveson N. G., (Eds.), *Resilience Engineering Concepts and Precepts*,Farnham, UK: Ashgate

Woods, D. D. (2006). Essential characteristics of resilience. In E. Hollnagel, D. D. Woods, & N. Leveson (Eds.), *Resilience Engineering: Concepts And Precepts* (pp. 19–30). Adelshot, UK: Ashgate.

Woods, D. D., & Wreathall, J. (2008). Stress-strain Plot as a Basis for Addressing System Resilience. In E. Hollnagel, C. P. Nemeth, & S. W. A. Dekker (Eds.), *Resilience Engineering Perspectives: Remaining Sensitive to the Possibility of Failure* (pp. 143–158). Adelshot, UK: Ashgate.

Woods, D., & Branlat, M. (2010). Hollnagel's test: being "in control" of highly interdependent multi-layered networked systems. *Cognition, Technology & Work*, 12(2), 95–101.

Woods, D. D., & Branlat, M. (2011). Basic Patterns in How Adaptive Systems Fail. In E. Hollnagel, J. Pariès, D. D. Woods, & J. Wreathall (Eds.), *Resilience Engineering in Practice* (pp. 127–144). Farnham, UK: Ashgate.

Wreathall, J., & Woods, D. D. (2008). The Stress-Strain Analogy of Organisational Resilience. In: Hollnagel, E. Nemeth, C. & Dekker, S.W.A. (Eds.) *Remaining Sensitive to the Possibility of Failure*. Burlington, VT, Ashgate Publishing Co.