Safety-II and the management of uncertainty

Context

Time to read: 

Foreword:

This paper discusses an approach to safety management and risk management that is called “Safety-II”. It is an approach to achieving good safety performance through focusing on, and learning from, “what works well” in the workplace, which includes avoiding hazards and, importantly, harnessing opportunities.

The paper is particularly relevant to industries in which inherent risks to operations present major potential safety consequences – which includes opportunities to achieve significant positive outcomes, not just the potential for significant negative outcomes that are perhaps what people first think of when they think of managing risks to achieving good safety performance. Examples of such industries include (but are not limited to) aviation, construction, energy and power, engineering, infrastructure, heavy industry and manufacturing, mining and oil & gas.

The paper focuses on safety whilst noting that safety is often combined with occupational health, mental health and wellbeing, and / or environmental management in businesses and their management systems.

Introduction

This paper asks how safety management and risk management can be used together in today’s modern and dynamic workplaces to help us achieve the best safe outcomes for organisations and their stakeholders.

We address and discuss the following:
1. “Safety risk” – should we reposition our understanding of what it actually is?
2. How do we define safety?
3. Is safety compliance-driven, or risk-led?
4. How does Safety-II link to the management of risk?
5. Should you change how you look at safety?

The paper concludes with some closing thoughts for readers to consider.

Risk, and the management of risk

We hear the term “safety risk” all the time. By saying these words, we are implying that “*there is a safety risk to achieving activity ABC*”.

Rather than saying “safety risk”, is there is a better way to describe how safety management and the management of risk can work together, so that activities achieve safe outcomes?

The term “risk” is defined in the international standard for risk management, ISO 31000:2018 (Risk Management – Guidelines), as “the effect of uncertainty on objectives”. A risk can be an opportunity or a threat. The potential impacts or consequences of a risk, therefore, including those of a safety nature, can be good or bad.

Figure 1 – risk – uncertainty – different impacts and outcomes

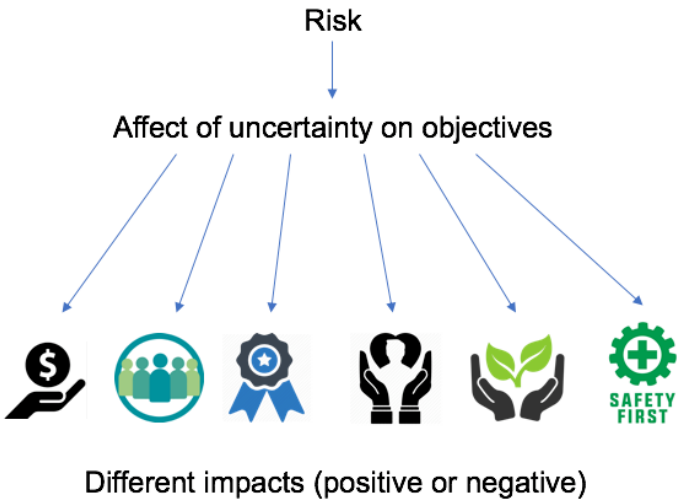

When we identify and assess our risks (opportunities and threats), we should use a good risk management process to identify, assess, manage and review all risks in the context of our situation, and the objectives we are seeking to achieve. In doing so, we should be ensuring that safety impacts, or consequences, are considered together with all other impacts / consequences, such as those of a financial, reputational or operational nature. This includes seizing opportunities as well as managing threats. As we manage our risks and review whether our actions are having the desired effect, we need to see how this is helping us to achieve our objectives.

Figure 2 – a 4-step risk management process (adhering to ISO 31000)

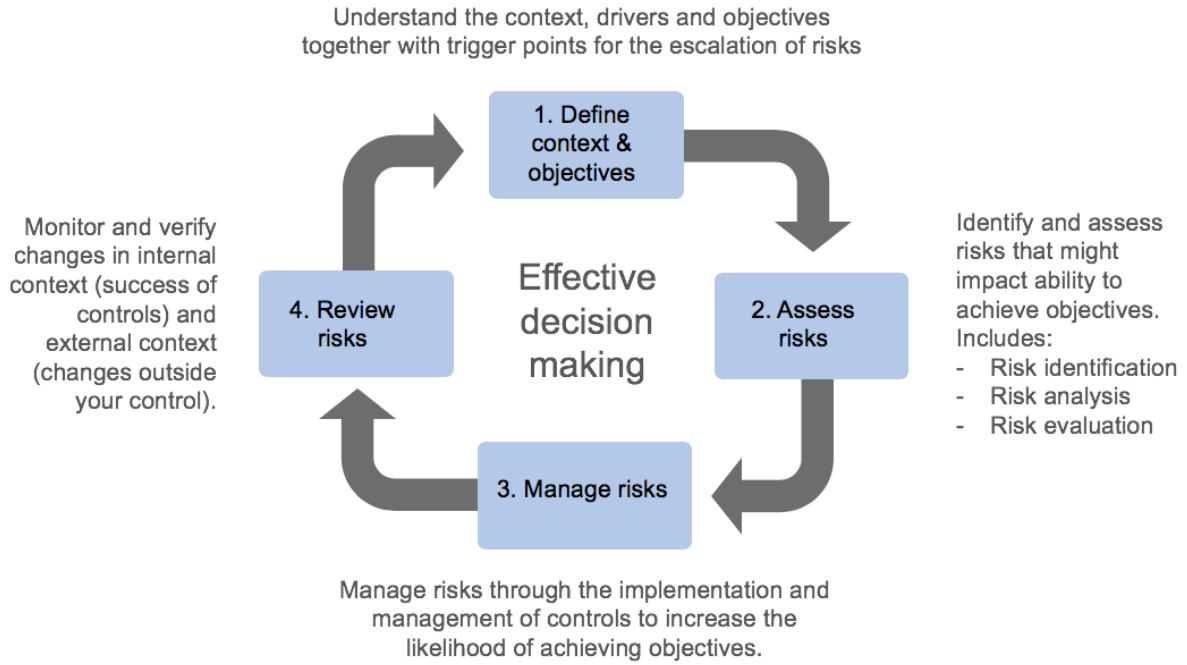

Achieving work safely

Achieving work safely is an outcome of performing work properly, which includes managing risk well. To manage risk well, we design and implement the right controls, and ensure the right knowledge and care is taken to perform the activity, with the right amount of time to perform it whilst ensuring we understand the context of the situational environment at all times. When we assess, manage and review risks to our activities, safety is part of an integrated set of considerations that we consider.

How do we define safety?

Interestingly, there is no international consensus towards a common definition of the word “safety”. The origin of the word seems to come from the Old French word, “sauf”, meaning “uninjured” or “unharmed”. [1]

A Google Search on “safety definition” returns a large number of suggestions, but there is no common definition from leading institutes around the world, such as the ISO, OSHA, the HSE in the UK, Safe Work Australia, the ANSI, the BSI or Standards Australia. Indeed, it’s not easy to find a definition for safety on the websites of these organisations.

The absence of a common definition for safety is not a hindrance. We all assume we know what safety means. Perhaps it is “obvious”; maybe that is why international standards and industry bodies do not seek to provide a common definition.

When we do find definitions, they usually focus on “avoiding negative outcomes”.

For example, the American National Standards Institute defines safety as: *the freedom from unacceptable risk*. [2] Unacceptable risk in this context refers to a risk for which the probability of occurrence is too high.

The International Standards Organisation, in ISO 45001:2018 (*Occupational health and safety management systems — Requirements with guidance for use*), refer to OH&S risk but not a direct definition for safety. Their definition of OH&S risk is based on negative risks. It is described as: *combination of the likelihood of occurrence of a work-related hazardous event(s) or exposure(s) and the severity of injury and ill health that can be caused by the event(s) or exposure(s).* [3]

We certainly want to avoid negative safety consequences in the work we perform. However, are these definitions too negative, and do they give rise to a mindset and approach that ignores and fails to learn from well performed work?

Taking a positive approach to safety…

Instead of focusing people's minds on avoiding unwanted outcomes and overcoming negative, risks events and incidents, is there a benefit to defining safety in terms of what should go well, and to focus on achieving positive consequences and good outcomes? Related to this, can safety be integrated into the "reliability and resilience of our operations", and as part of this, linked to quality management?

In 1987, in an article for the California Management Review, professor Karl Weick (who has since written extensively about risk and resilience) introduced the idea of reliability a dynamic non-event [4]:

> Reliability is dynamic in a sense that it is an ongoing condition in which problems are momentarily under control due to compensating changes in components. Reliability is invisible in at least two ways. First, people often don't know how many mistakes they could have made but didn't, which means they have at best only a crude idea of what produces reliability […] Reliability is also invisible in the sense that reliable outcomes are constant, which means there is nothing to pay attention to.

There is direct linkage with the concept of reliability by Weick to the work of Professor Erik Hollnagel. Hollnagel defines being safe as: *the outcome of what is being done is as we expect – that the activities we undertake will be successful*. [5] This is a positive definition. If we apply this this definition and way of thinking, can we harness the power of focusing on, and learning from "what goes well"?

Is safety compliance-driven, or risk-led?

As we have seen with common definitions for safety, much of the traditional focus in safety circles is about elimination, or reduction, of unwanted incidents and events. [6] The traditional, or conventional, approach to safety is defined as "the pursuit of a condition where nothing goes wrong", with a need to ensure compliance measures are in place to stop things going wrong. This has been termed by Hollnagel as ***'Safety-I'*** *thinking*. [7] From this perspective, the purpose of safety management is to implement measures to ensure the number of incidents and events is kept as low as possible, or as low as is reasonably practicable. In this perspective, the term "safety consequence" (which we discussed earlier can be positive or negative) in Safety-I is always negative.

Is quantifying safety by *what goes wrong* the *right* approach…?

It is natural to want to fix things that go wrong, and to fix faults that are found in safety management, but this focus should not be at the expense of looking at what goes right, and how to ensure things go right on a continuous basis.

As Hollnagel describes, when we quantify safety by measuring what goes wrong rather than what goes well, we are implying that the safer an activity or a system is, the less need there is to measure it.

Therefore, if a system is absolutely safe, in this approach we have, theoretically, nothing to measure. In control theory, this is known as "*the regulator paradox*" (and it applies to other aspects other than safety, such as managing financial risk). [8]

The regulator paradox is this: if something rarely goes wrong, perhaps never has till now, it is almost impossible to know how well it is working if we just look at statistics of "what goes wrong". This has been described as follows: [9]

> The task of a regulator is to eliminate variation, but this variation is the ultimate source of information about the quality of its work. The better the job the regulator does, the less information it gets about how to improve.

A compliance-driven approach which looks at avoiding negative outcomes is what Hollnagel terms **Safety-I**. Safety-I defines safety as *a condition where the number of adverse outcomes (accidents / incidents / near misses) is as low as possible.* [10]

Safety-I achieves its goals by either finding out what went wrong and fixing it, or preventing a change or deviation from "normal" by implementing barriers (be they physical, digital or temporal). In order for it to work, processes must be familiar to us so that we can detect potential problems quickly, and stop them from happening before they become real. It is largely a "directive" approach, with limited engagement with, and input from, people who actually undertake the work.

It is fairly commonplace to see "safety dashboards" in businesses that are choc-full of statistics of ratios that look at what's gone wrong (e.g. with a variety of "frequency ratios" such as TRIFR, LTIFR, CIFR and other lag indicators).

In publications and at conferences, we tend to hear about case studies and examples of what has gone wrong and what we should learn as a result. There is definitely value in learning from mistakes, but this should not be at the detriment of learning from “what we do well”.

Consider the reporting of safety through incidents, frequency ratios, and “event probability of occurrence” as a way to measure safety performance. If we have determined with the aid of data a statistical probability of an event occurring as a “one in every 10,000 occurrences”, this means we should expect something to go wrong in one event every 10,000 times of the operation. What does monitoring against this metric actually tell us? Are we waiting for the one in 10,000 events to occur, and to react to it when it does? Would we be better-informed and able to proactively act on risks that exist – both opportunities and threats – if we focused on knowing and understanding all the good, small adjustments to our operation which ensured a positive outcome thousands of times? [11]

Risk management for Safety-I

Risk management for Safety-I is limited in what it can achieve, because it focuses on known routes and causes of failure, it does not “cast the net of possibilities” any wider. It focuses on threats, not opportunities. Risk reviews of operations may be performed just once, at the start of an activity, or perhaps at occasional points in time, including in reaction to negative events occurring. It may well be performed in a thorough manner as part of a compliance process, but in formal application at least, it is not dynamic or proactive, or self-adjusting. It does not instil a risk-informed mindset of spotting dynamic changes to assess and act upon, including opportunities. [12]

A Safety-I approach is always on the defensive, reacting to changing circumstances rather than anticipating and adapting in advance.

As we have mentioned earlier, avoiding negative safety events, and taking steps to avoid known negative risks is important. However, it should not dominate our attention. If we are to show true commitment, care and trust in the people who undertake the work that needs to get done, we should focus on finding out *what goes right*. This means focusing on what we normally pay little attention to, and learning from it in a Plan-Do-Check-Act type of closed loop. Shining a spotlight on “what goes right” requires us to look at things differently and it is an approach that good, proactive risk management can support.

Safety-II – focusing on "what goes well"

Safety-II as defined by Hollnagel is a different mindset to Safety-I. It is about focusing on what goes well. This includes engaging fully and properly with our workforce, because they have a crucial role to play in this. By looking at outcomes as "wanted events", we view safety consequences arising from risks as positive, as opportunities to succeed. As part of this, we need to look at the constantly changing dynamic factors to activities people perform, and how people adjust for such changes. This is why properly engaging with the people who undertake the physical activities, and ensuring we learn from them, is so valuable and important.

In line with the principles of resilience, Safety-II can be defined as: *the ability to succeed under both expected and unexpected conditions, so that the number of intended and acceptable outcomes is as high as possible.* [13]

Safety-II and complexity in today's world

There is an inherent link with Safety-II and systems theory, and how we deal with complexity and resilience. Whilst complexity and resilience is not the subject of this paper, it is important to note that these concepts are inherent in Safety-II.

Complexity and resilience are important matters for organisations to address today. Organisations exist in complex eco-systems where entities (e.g. businesses, regulators, individuals or teams) are interrelated in dynamic, uncertain, and often ambiguous environments. Aspects of this system, such as levels of detail, interactions between parties and probabilities of events (negative or positive) occurring are not universal constants, they change, and often change quickly.

For groups and parties of people working within these complex systems, their efforts to work with each other give rise to impacts that can create the emergence of a new property in the system. Adjustments are often required to manage these new properties. This is increasingly common today, and will continue to be so in future, with the advent of Industry 4.0, and with devices and controls being increasingly automated and connected to each other in a "smart eco-system".

Managing safety and risk in today's complex world

Safety-II acknowledges that whilst the moving parts in a system are not perfectly predictable (i.e. there is some inherent uncertainty with undertaking them), we usually have some, and often considerable, knowledge about them. Only on rare occasions do we have no knowledge about them at all. Minor adjustments to activities are constantly required to adapt to changes, large or small, in response to a dynamically changing environment.

This context of "adaptive work" is a key aspect of Safety-II. Safety-II helps us to respond to a world that is complex and has many variable factors.

To ensure our management systems respond to complexity, we need to make sure they incorporate the need to make performance adjustments and adapt to variability, and to learn from such occurrences in order to continually achieve positive outcomes. They should not be unwieldy and inflexible by being too rigid. Rigidity does not work in resilient systems. [14]

Linked to focusing on "what goes well", safety should be woven into the fabric of an integrated management system; it should not exist separately or standalone. As an example of an integrated approach, consider a modern supply chain. It is likely to be more complex and interlinked today than it was in the past, with operations upstream coupled to operations downstream in sophisticated ways and probably involving many parties. The close connections we create in this supply chain create a complex environment (system) and a resulting need to cater for regular adjustments to ensure resilience. Harnessing the right adaptability to adjust can help us to achieve reliably safe outcomes in all supply chain activities.

Human factors of course play an important part in Safety-II, as part of the constant adjustment and adaptation that we recognise is required. Ingrained into the Safety-II ethos is that human performance, individually and collectively, is always variable. [15] Safety-II explicitly takes into account that systems function because people have the skills and awareness to make adjustments to suit the nature of specific conditions they face at specific points in time. [16] Regular performance adjustments that people make are seen as a positive form of variability, allowing the system to flex to suit a specific situational circumstance. We value people's skills, we do not seek to constrain them. Learning from the good adjustments people make by engaging properly with them can help us maintain a positive focus on achieving good safe outcomes along with all other intended outcomes (including financial targets and everything else).

How does Safety-II link with the management of risk?

Let's return to our earlier point, that good risk management is about recognising uncertainty, identifying it within our context, and working out what we should do about it to achieve our objectives. Managing risk is not solely about managing negative possible outcomes, which many people assume it to be. It is about realising potential opportunities to achieve objectives and success. Many organisations differentiate the "good and the bad" as opportunities (good) and risks (bad). This is perfectly fine; the main thing is that "the positives" are focused on as much as "the negatives".

Risk management is not an occasional spot-check, or an activity risk review held at a point in time, or infrequently at occasional "check point" stage gates (valid though such gate reviews are). Good risk management is a continual process that supports decision-making, and the making of good decisions in the face of uncertainty. The continual review to understand our activities and context as things change in our complex systems, and looking for opportunities to achieve objectives, ties Safety-II and risk management together.

Safety-II involves identifying, assessing and responding to change *before* anything happens. These adjustments should be for outcomes that are "wanted" (i.e. opportunities) as well as outcomes to be avoided (i.e. threats). The ethos of Safety-II is that people adjust their work to ensure that the outcomes are what we want - we are proactive about it to ensure it happens. It is not enough to be reactive. This ethos can be aided by proactive and integrated risk management. Through integrated risk management we can identify insights into possible patterns and adjust to changing conditions to keep things working well. We measure, analyse, learn and act upon what occurs. By doing so, we manage a “whole activity”, with safety integrated into all considerations as part of focusing on our overall objectives.

So, what’s our appetite for risk?

We have to take risks to succeed with our ventures and plans. Some risks should not be taken, and others can and should be. Defining our risk appetite and tolerance for different types of risk is a good way to provide clarity to the risks we are willing to take (opportunities) and those we are willing to accept in order to achieve our objectives, and where our thresholds lie for them. In a complex system, having clarity on our risk appetite and tolerance can provide us with guidelines on which dynamic changes are appropriate to keep us within our appetite, or if necessary, within our tolerance.

Our appetite and tolerance for risk can be expressed in various ways. Perhaps the easiest way is to write detailed procedures and rules on what you can and can’t do. Procedures and rules are important, but in a complex environment they must be practical, and support people to work in a dynamic environment, with the right degree of flexibility to adapt (just as complex natural ecosystems do).

Appetite and tolerance should take into account integrated aspects. For example, financial appetite and tolerance cannot, and should not, be decoupled from human-based aspects and incentives – which includes safety appetite and tolerance. Businesses have to make commercial decisions every day, some large, some small, and these decisions should take all aspects of appetite and tolerance into account.

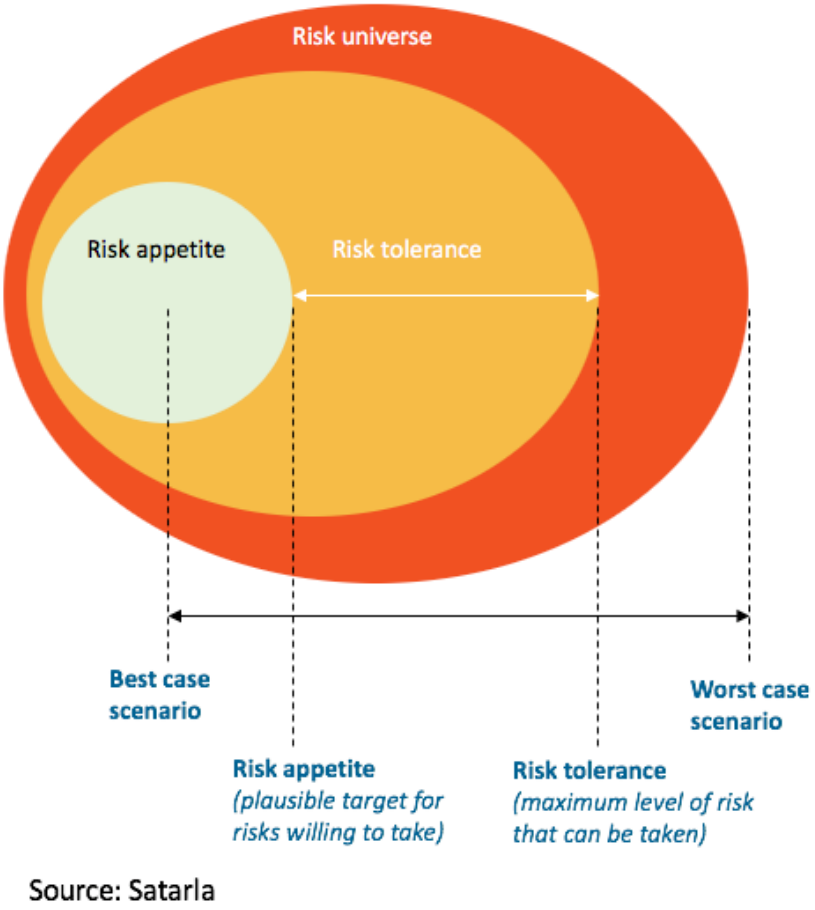

Figure 3 – risk appetite and tolerance

Good risk management contributes to safety being part of "good work"

As we discussed earlier, an important aspect of Safety-II is that the complexity of activities and interactions required to deliver outcomes safely means that we work in a system that is subject to change and a degree of unpredictability. We are often required to adapt tasks that we have designed (Work-As-Imagined) to get the job done (Work-As-Done). How we recognise and make such adjustments and remain resilient is part of Safety-II. This is both to strive to ensure opportunities become "wanted events", and to prevent threats turning into "unwanted events".

Key to this is a need to give people autonomy in decision-making, to minimise bureaucracy, and to ensure everyone seeks out the views of others for diversity of thought. Autonomy, minimal bureaucracy and diversity of thought creates trust, and a feeling of being trusted to make a decision.

Establishing this ethos requires rigorous assurance up front to provide people with a platform for success. When people have a clearly defined risk appetite and tolerance, they need to be trusted to take and manage risks, and escalate matters when they feel they need to, within a "psychologically safe" environment. This requires strong leadership and a good culture to be in place.

By starting with a focus at the front-end of designing and procuring work activities well, and continuing to maintain this focus on well-managed work in execution and delivery, strong performance and delivery of financial and all other goals can be coupled with achieving strong safety performance.

To assist the application of Safety-II, risk management can provide people with tools and techniques to think through a range of outcomes, including thinking through 'what's the worst that could happen?' and also 'what's the best that could happen?' with credible scenarios reviewed. The Pre-mortem (negative) and Pro-mortem (positive) scenario analysis technique is an example of a risk technique we can use to address our best and worst case scenarios.

Should you make a change to how you look at safety?

Are you currently using a compliance-driven Safety-I as your approach to safety, and is it working?

Would the application of a Safety-II approach for safety management, which focuses on "what goes well" and leverages good risk management, add value to your organisation?

How the Safety, Risk and Quality functions can work together

We have seen that when safety is being carried out at physical operations, the need to dynamically assess the situation and context every day is vital. In order for our operations to be successful, they have to be set up for success through the planning and design of our activities, which requires good risk and safety management to be central to activity planning.

Risk and Safety advisors need to work together with their operations teams to facilitate good outcomes. For example, when a risk review is being held, always consider the "human factors" alongside financial / commercial factors.

Safety can also benefit from being tied to a Quality Management System (QMS). A QMS typically stretches across an organisation as an overarching "way of doing things". Practices such as Six Sigma focus on repeatable quality outcomes – much the same as the Safety-II approach that we have written about.

Conclusion

This paper has looked at traditional "Safety-I" compliance-orientated ways of managing safety and compared it with a different "Safety-II" approach of focusing on "what goes well".

To be sure, there are important aspects of the compliance-driven Safety-I approach that remain valid, to learn from mistakes and problems that occur.

By applying a Safety-II approach to focus on what works well, and the regular and consistent achievement of successful outcomes, we can be positively-minded about achieving successful safe outcomes from our activities. Good risk management can help us to achieve this, especially in the dynamic modern systems that we operate, where continual adjustments are required in order to maintain resilient operations.

For policies, management systems and standards, Safety-II can be used to focus on and to encourage positive learning from what goes right. It can help us to ensure that documentation is at the appropriate level of detail; that it is not cumbersome and that it does not try to cover every possible situation in which unwanted events can occur (which in reality is not possible).

By focusing on what goes well in our activities, by trusting people to use their skills and awareness to achieve “wanted events”, and by using clearly defined risk appetite and tolerance to help people make decisions, we can make the best risk-informed decisions in order to achieve safe outcomes and to achieve our overall objectives.

About the Author

Gareth is an international Risk consultant working across various industries and sectors, particularly those of a capital and asset-intensive nature.

The author wishes to thank to Professor Erik Hollnagel for his advice and guidance. Professor Hollnagel is CEO at the Institute of Resilient Systems +, and a Visiting Professor at many international universities. He is the author of the book, *Safety-I and Safety-II - The Past and Future of Safety Management*.

References

1 Hollnagel, E, 2014, *Safety-I and Safety-II*, CRC Press, p.1
2 Hollnagel, E, 2014, *Safety-I and Safety-II*, CRC Press, p.50
3 ISO 45001:2018 definitions
4 Weick, KE, 1987, *Organizational Culture as a Source of High Reliability*, California Management Review, Vol 29, pp 112-128 (p. 118)
5 Hollnagel, E, 2014, *Safety-I and Safety-II*, CRC Press, p.??
6 Hollnagel, E, 2014, *Safety-I and Safety-II*, CRC Press, p.49
7 Hollnagel, E, 2014, *Safety-I and Safety-II*, CRC Press, p.50
8 Hollnagel, E, 2014, *Safety-I and Safety-II*, CRC Press, p.50
9 Hollnagel, E, 2014, *Safety-I and Safety-II*, CRC Press, p.11
10 Hollnagel, E, 2014, *Safety-I and Safety-II*, CRC Press, p.49
11 Hollnagel, E, 2014, *Safety-I and Safety-II*, CRC Press, p.46
12 Hollnagel, E, 2014, *Safety-I and Safety-II*, CRC Press, p.57
13 Hollnagel, E, 2014, *Safety-I and Safety-II*, CRC Press, p.134
14 Hollnagel, E, 2014, *Safety-I and Safety-II*, CRC Press, p.113
15 Hollnagel, E, 2014, *Safety-I and Safety-II*, CRC Press, p.127
16 Hollnagel, E, 2014, *Safety-I and Safety-II*, CRC Press, p.137