

A tale of two safeties

HOLLNAGEL Erik^{1,2}

1 Professor, Institute of Regional Health Research, University of Southern Denmark

2. Chief Consultant, Center for Quality, Region of Southern Denmark, P. V. Tuxens Vej 5, DK-5500 Middelfart, Denmark
(erik.hollnagel@rsyd.dk)

Abstract: The sustained existence of modern societies depends on the safe and efficient functioning of multiple systems, functions, and specialised services. Because these often are tightly coupled, safety cannot be managed simply by responding whenever something goes wrong. Both theory and practice make clear that safety management that follows developments rather than leads them runs a significant risk of lagging behind and of becoming reduced to uncoordinated and fragmentary fire-fighting. (The same, of course, goes for the management of quality and productivity.) In order to prevent this from happening, safety management must look ahead, not only to avoid that things go wrong but also – and more importantly – to ensure that they go right. Proactive safety management must focus on how everyday performance usually goes well rather than on why it occasionally fails, and must actively try to improve the former rather than simply prevent the latter.

Keyword: Safety-I; safety-II; resilience engineering; performance variability; successes

1 Safety as the freedom from unacceptable risk

Safety has traditionally been defined as a condition where nothing goes wrong. Or rather, since we know that it is impossible to ensure that nothing goes wrong, as a condition where the number of things that go wrong is acceptably small (See in the part of Appendix, *1). This is, however, an indirect and somewhat paradoxical definition since safety is defined by its opposite, by what happens when it is missing. As a consequence of this definition, safety is also measured indirectly, not by its presence or as a quality in itself, but by the consequences of its absence.

In relation to human activity it makes good practical sense to focus on situations where things go wrong, both because such situations by definition are unexpected and because they may lead to unintended and unwanted harm or loss of life and property. An early example is the collapse of the Rialto Bridge in Venice, when it became overloaded with spectators at the wedding of the Marquess of Ferrara in 1444. (Many spectacular accidents have, of course, happened before that, but the historical record is sketchy and incomplete.) The bridge collapse is characteristic of the classical safety concerns, which addressed risks related to passive technology and structures such as buildings, bridges, ships, *etc.* This concern was reinforced by the needs of the second

industrial revolution, around 1750, which was marked by the invention of a usable steam engine. The rapid mechanisation of work that followed led to a growing number of hitherto unknown types of accidents, where the common factor was the breakdown, failure, or malfunctioning of active technology. Andrew Hale and Jan Hovden^[1] have characterised this as the age of technology, in which safety concerns focused on guarding machinery, stopping explosions and preventing structures from collapsing. The focus on technology as the main – or even only – source of both problems and solutions in safety was successfully maintained until 1979, when the accident at the Three Mile Island nuclear power plant demonstrated that safeguarding technology was not enough (See in the part of Appendix, *2). The TMI accident brought to the fore the role of human factors – or even of the human factor – and made it necessary to consider human failure and malfunctioning as a potential risk. Seven years later the loss of the space shuttle Challenger, reinforced by the accident in Chernobyl, required yet another extension, this time by adding the influence of organisational failures and safety culture to the common lore.

Throughout the ages, the starting point for safety concerns has been the occurrence, potential or actual, of some kind of adverse outcome, whether it has been categorised as a risk, a hazard, a near miss, an incident, or an accident. Historically speaking, new

Received date: February 28, 2013

types of accidents have been accounted for by introducing new types of causes (*e.g.*, metal fatigue, ‘human error’, organisational failure) rather than by challenging or changing the basic underlying assumption of causality. We have therefore through centuries become so accustomed to explaining accidents in terms of cause-effect relations – simple or compound – that we no longer notice it. And we cling tenaciously to this tradition, although it has become increasingly difficult to reconcile with reality.

1.1 Habituation

An unintended but unavoidable consequence of associating safety with things that go wrong is a creeping lack of attention to things that go right. The psychological explanation for that is called habituation, a form of adaptive behaviour that can be described as non-associative learning. Through habituation we learn to disregard things that happen regularly, simply because they happen regularly. The formal definition of habituation is a “response decrement as a result of repeated stimulation”^[2]. In academic psychology, habituation has been studied at the level of neuropsychology and also usually been explained at that level^[3].

It is, however, entirely possible also to speak about habituation at the level of everyday human behaviour – actions and responses. This was noted as far back as in 1890, when William James, one of the founding fathers of psychology, wrote that “habit diminishes the conscious attention with which our acts are performed.”^[4] In today’s language it means that we stop paying attention to something as soon as we get used to doing it. After some time we neither notice that which goes smoothly, nor do we think it is necessary to do so. This applies both to actions and their outcomes – both what we do ourselves and what others do.

From an evolutionary perspective, as well as from the point of view of an efficiency-thoroughness trade-off^[5], habituation makes a lot of sense. While there are good reasons to pay attention to the unexpected and the unusual, it may be a waste of time and effort to pay much attention to that which is common or similar. To quote James^[4] again: “Habitual actions are certain, and being in no danger of going astray from their end, need no extraneous help” (p.

149). Reduced attention is precisely what happens when actions regularly produce the intended and expected results and when things ‘simply’ work. When things go right there is first of all no difference between the expected and the actual, hence nothing that attracts attention or initiates an arousal reaction. Neither is there any motivation to try to understand why things went well: they obviously went well because the system – people and technology – worked as it should and because nothing untoward happened. While the first argument – the lack of a noticeable difference between outcomes – is acceptable, the second argument is fatally flawed. The reason for that will become clear in the following.

2. Looking at what goes wrong rather than looking at what goes right

To illustrate the consequences of looking at what goes wrong rather than looking at what goes right, consider Fig. 1. This represents the case where the (statistical) probability of a failure is 1 out of 10,000 – technically written as $p = 10^{-4}$. This means that for every time we expect that something will go wrong (the thin line), there are 9,999 times where we should expect that things will go right and lead to the outcome we want (the grey area). The ratio of 1:10,000 corresponds to a system or organisation where the emphasis is on performance^[6]; the ratio would be even more extreme for an ultrasafe system.

As an example of this, consider the train collision in Buizingen, Belgium on 15 February 2010^[7]. Two trains, carrying 250–300 people, collided in snowy conditions during the morning rush hour. The trains apparently collided “laterally” at a set of points at the exit of Halle station. Eighteen people were killed and 162 injured, and there was major damage to the tracks. The investigation found that one of the trains had passed a red signal without stopping (SPAD or Signal Passed At Danger), and that this could be a contributing cause to the collision. On further investigation, it was found that there were 130 SPAD events in Belgium in 2012, of which one third were serious. But it was also estimated that there were about 13.000.000 cases of trains stopping at a red signal. The probability of a SPAD was therefore 10^{-5} , the probability of a serious SPAD was $3.3 \cdot 10^{-6}$, and the probability of the accident was $7.7 \cdot 10^{-8}$.

Another example can be found in the statistics for

Frankfurt airport. In 2011 there were a total of 490,007 movements, but only 10 infringements of separation and 11 runway incursions. This corresponds to a ratio of $2.04 \cdot 10^{-5}$ and $2.25 \cdot 10^{-5}$, respectively, or roughly 2 cases out of every 100.000.




Fig.1 The imbalance between things that go right and things that go wrong.

The tendency to focus on what goes wrong is reinforced in many ways. It is often required by regulators and authorities; it is supported by models and methods; it is documented in countless databases and illustrated by almost as many graphs; it is described in literally thousands of papers, books, and conference proceedings; and there are an untold number of experts, consultants, and companies that constantly remind us of the need to avoid risks, failures, and accidents – and of how their services can help to do just that. The net result is abundant information both about how things go wrong and about what must be done to prevent this from happening. The focus on failures also conforms to our stereotypical understanding of what safety is and on how safety should be managed, cf., above. The recipe is the simple principle known as ‘find and fix’: look for failures and malfunctions, try to find their causes, and try to eliminate causes and/or improve barriers.

One unfortunate and counterproductive consequence of this is that safety and core business (production) compete for resources; this means that investments in safety are seen as costs, and therefore (sometimes) hard to justify or sustain. Another consequence is that learning is limited to that which has gone wrong, which means that it only happens infrequently and only uses a fraction of the data available.

The situation is quite different when it comes to that which goes right, *i.e.*, the 9,999 events out of the 10,000. A focus on what goes right receives little

encouragement. There is no demand from authorities and regulators to look at what works well, and if someone should want to do so there is little help to be found; we have few theories or models about how human and organisational performance succeeds, and few methods to help us study how it happens; examples are few and far between^[8], and actual data are difficult to locate; it is hard to find papers, books or other forms of scientific literature about it; and there are few people who claim expertise in this area or even consider it worthwhile. It furthermore clashes with the traditional focus on failures, and even those who find it a reasonable endeavour are at a loss when it comes to the practicalities: there are no simple methods or tools and very few good examples to learn from.

Yet one interesting consequence of this perspective is that safety and core business no longer compete for resources; what benefits one will also benefit the other. Another consequence is that learning can focus on that which has gone right, which means that there are literally countless opportunities for learning, and that data are readily available – once the attention is turned away from failures.

3 Safety-I: Avoiding that things go wrong

The traditional definition of safety as a condition where the number of adverse outcomes (accidents / incidents / near misses) is as low as possible can be called Safety-I. The purpose of managing Safety-I is consequently to achieve and maintain that state. The U.S. Agency for Healthcare Research and Quality, for instance, defines safety as the “freedom from accidental injury,” while the International Civil Aviation Organization defines safety as “the state in which 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 risk management.”

The ‘philosophy’ of Safety-I is illustrated by Fig. 2. Safety-I promotes a bimodal or binary view of work and activities, according to which they either succeed or fail. (This is of course in good agreement with the standard methods for representing accidents and risks, which all are based on some form of branching tree.) When everything works as it should (‘normal’ functioning), the outcomes will be acceptable; things go right, in the sense that the number of adverse

events is acceptable small. But when something goes wrong, when there is a malfunction, human or otherwise, this will lead to a failure (an unacceptable outcome). The issue is therefore how the transition from normal to abnormal (or malfunction) takes place, *e.g.*, whether it happens through an abrupt or sudden transition or through a gradual ‘drift into failure’. According to the logic of Safety-I, safety and efficiency can be achieved if this transition can be blocked.

The focus on failures creates a need to find the causes of what went wrong. When a cause has been found, the next logical step is either to eliminate it or to disable suspected cause-effect links. Following that, the outcome should then be measured by counting how many fewer things go wrong after the intervention. Safety-I thus implies what might be called a ‘hypothesis of different causes,’ namely that the causes or ‘mechanisms’ of adverse events are different from those of events that succeed. If that was not the case, the elimination of such causes and the neutralisation of such ‘mechanisms’ would also reduce the likelihood that things could go right, hence be counterproductive.

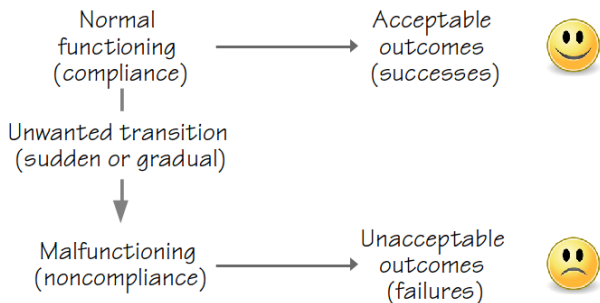


Fig.2 The Safety-I view of failures and successes.

Safety-I tacitly assumes that systems work because they are well designed and scrupulously maintained, because procedures are complete and correct, because designers can foresee and anticipate even minor contingencies, and because people behave as they are expected to – and more importantly as they have been taught or trained to do. This unavoidably leads to an emphasis on compliance in the way work is carried out.

The background for the Safety-I perspective is found in well-understood, well-tested, and well-behaved systems. It is characteristic of such systems that there is a high degree of reliability of equipment, that

workers and managers are vigilant in their testing, observations, procedures, training, and operations, that staff is well trained, that management is enlightened, and that good operating procedures are in place. If these assumptions are correct, humans – as ‘fallible machines’ – are clearly a liability and their performance variability can be seen as a threat. According to the logic of Safety-I, the goal – the coveted state of safety – can be achieved by constraining all kinds of performance variability. Examples of frequently used constraints are selection, strict training, barriers of various kinds, procedures, standardisation, rules, and regulations. The undue optimism in the efficacy of this solution has extended historical roots. But whereas the optimism may have been justified to some extent a hundred years ago, it is not so today. The main reason is that the work environment has changed dramatically, and to such an extent that the assumptions of yesteryear are no longer valid.

3.1 Safety-I: Reactive safety management

The nature of safety management clearly depends on the definition of safety. From a Safety-I perspective, the purpose of safety management is to make sure that the number of adverse outcomes is kept as low as possible – or as low as reasonably practicable^[9]. A good example of that is provided by the WHO research cycle shown in Fig. 3. The figure shows a repeated cycle of steps that begins when something has gone wrong so that someone has been harmed. In health care, ‘measuring harm’ means counting how many patients are harmed or killed and from what type of adverse events. In railways, accidents can be defined as “employee deaths, disabling injuries and minor injuries, per 200,000 hours worked by the employees of the railway company” or “train and grade crossing accidents that meet the reporting criteria, per million train miles”. Similar definitions can be found in every domain where safety is a concern.

This approach to safety management is reactive, because it based on responding to something that either has gone wrong or has been identified as a risk – as something that could go wrong. The response typically involves looking for ways to eliminate the cause – or causes – that have been found, or to control the risks, either by finding the causes and eliminating

them, or by improving options for detection and recovery. Reactive safety management embraces a causality credo, which goes as follows: (1) Adverse outcomes (accidents, incidents) happen when something goes wrong. (2) Adverse outcomes therefore have causes, which can be found and treated.

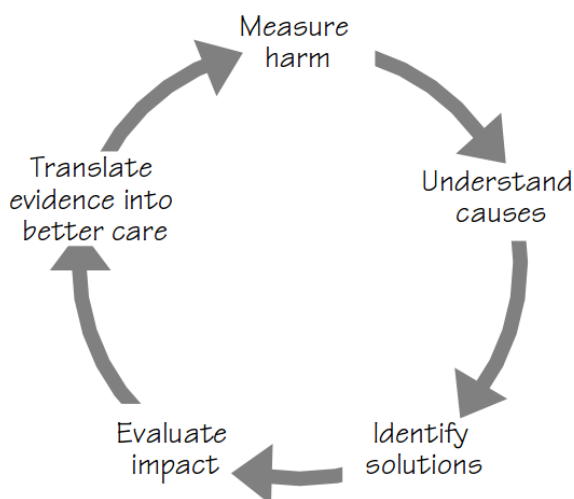


Fig.3 Reactive safety management cycle (WHO).

From a Safety-I perspective, the purpose of safety management is to keep the number of accidents and incidents as low as possible by reacting when an unacceptable event has occurred. Such reactive safety management can in principle work if events do not occur so often that it becomes difficult or impossible to take care of the actual work, *i.e.*, the primary activities. But if the frequency of adverse events increases, the need to respond will sooner or later require so much capacity that the reactions both become inadequate and partly will lag behind the process. In practice, it means that control of the situation is lost and with that the ability effectively to manage safety^[10]

Practical examples of this condition are easy to find. Severe weather – tornadoes or typhoons – may easily exhaust the capacity of the rescue services to respond. The same goes for forest fires or large oil spills – where the latter can come from ships or from the bottom of the sea. If patients are admitted to the emergency room at a rate that is higher than the rate by which they can be treated and discharged, the capacity to treat them will soon be exhausted. This can happen

during everyday conditions^[11], or during an epidemic^[12]. On a more mundane level, most industries (power plants, airlines, *etc.*) are struggling to keep ahead of a maelstrom of incident reports mandated by law. Even if only the most serious incidents are analysed, there may still be insufficient time to understand and respond to what happened.

Another condition is that the process being managed is familiar and sufficiently regular to allow responses to be prepared ahead of time (anticipation). The worst situation is clearly when something completely unknown happens, since time and resources then must be spent to find out what it is and work out what to do, before a response can actually be given. In order for reactive safety management to be effective, it must be possible to recognise events so quickly that the organisation can initiate a prepared response with minimal delay. The downside of this is that hasty and careless recognition may lead to inappropriate and ineffective responses.

4 Safety-II: Ensuring that things go right

As technical and socio-technical systems have continued to develop, not least due to the allure of ever more powerful information technology, systems and work environments have gradually become more intractable^[13]. Since the models and methods of Safety-I assume that systems are tractable, in the sense that they are well-understood and well-behaved, Safety-I models and methods are less and less able to deliver the required and coveted ‘state of safety.’ Because this inability cannot be overcome by ‘stretching’ the tools of Safety-I even further, it makes sense to consider whether the problem may lie in the definition of safety. One option is therefore to change the definition and to focus on what goes right rather than on what goes wrong (as suggested by Fig. 1). Doing so will change the definition of safety from ‘avoiding that something goes wrong’ to ‘ensuring that everything goes right’ – or more precisely to the ability to succeed under varying conditions, so that the number of intended and acceptable outcomes (in other words, everyday activities) is as high as possible. The consequence of this definition is that the basis for safety and safety management now becomes an understanding why things go right, which means an understanding of everyday activities.

Safety-II explicitly assumes that systems work because people are able to adjust what they do to match the conditions of work. People learn to identify and overcome design flaws and functional glitches, because they can recognise the actual demands and adjust their performance accordingly, and because they interpret and apply procedures to match the conditions. People can also detect and correct when something goes wrong or when it is about to go wrong, hence intervene before the situation becomes seriously worsened. The result of that is performance variability, not in the negative sense where variability is seen as a deviation from some norm or standard, but in the positive sense that variability represents the adjustments that are the basis for safety and productivity (Fig. 4).

In contrast to Safety-I, Safety-II acknowledges that systems are incompletely understood, that descriptions can be complicated, and that changes are frequent and irregular rather than infrequent and regular. Safety-II, in other words, acknowledges that systems are intractable rather than tractable^[13]. While the reliability of technology and equipment in such systems may be high, workers and managers frequently trade-off thoroughness for efficiency, the competence of staff varies and may be inconsistent or incompatible, and reliable operating procedures are scarce. Under these conditions humans are clearly an asset rather than a liability and their ability to adjust what they do to the conditions is a strength rather than a threat.

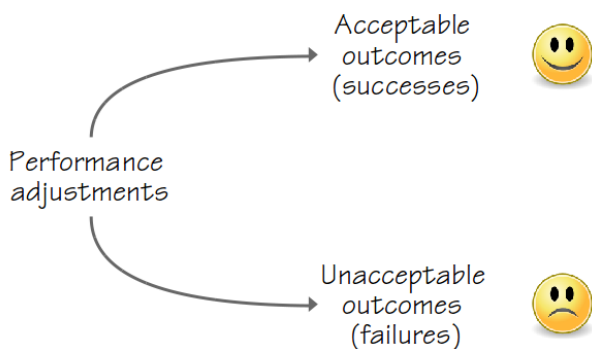


Fig.4 The Safety-II view of failures and successes.

Performance variability or performance adjustments are a *sine qua non* for the functioning of socio-technical systems, unless they are extremely

simple. Unacceptable outcomes or failures can therefore not be prevented by eliminating or constraining performance variability since that would also affect the desired acceptable outcomes. Instead efforts are needed to support the necessary improvisations and performance adjustments by clearly representing resources and constraints of a situation and by making it easier to anticipate the consequences of actions. Performance variability should be managed by dampening it if it is going in the wrong direction and amplifying it if it is going in the right direction. In order to do so it is necessary first to acknowledge the presence – and inevitability – of performance variability, second to monitor it, and third to control it. That is the remit of safety management according to Safety-II.

4.1 Safety-II: Proactive safety management

Safety-II management and resilience engineering both assume that everything basically happens in the same way, regardless of the outcome. This means that there is no need to have one set of causes and ‘mechanisms’ for things that go wrong (accident and incidents), and another for things that go right (everyday work). The purpose of safety management is to ensure latter, but by doing so it will also reduce the former. Although Safety-I and Safety-II both lead to a reduction in unwanted outcomes, they use fundamentally different approaches with important consequences for how the process is managed and measured – as well as for productivity and quality.

From a Safety-II perspective, the purpose of safety management is to ensure that as much as possible goes right, in the sense that everyday work achieves its stated purposes. This cannot be done by responding alone, since that will only correct what has happened. Safety management must instead be proactive, so that adjustments are made before something happens and therefore affect how it happens or even prevent something from happening. A main advantage is that early responses, on the whole, require a smaller effort because the consequences of the event will have had less time to develop and spread. And early responses can obviously save valuable time.

For proactive safety management to work, it is necessary to foresee what could happen with acceptable certainty and to have the appropriate means (people and resources) to do something about it. That

in turn requires an understanding of how the system works, of how its environment develops and changes, and of how functions may depend on and affect each other. This understanding can be developed by looking for patterns and relations across events rather than for causes of individual events. To see and find those patterns, it is necessary to take time to understand what happens rather than spend all resources on fire-fighting.

A trivial example is to ‘batten down the hatches’ when bad weather is approaching. While this expression has its origin in the navy, many people living on land – or on an oil rig – have also learned the value of preparing for a storm. In the financial world, proactive safety management is *de rigueur*; a financial institution that can only react will soon be out of business. In a different domain, the precautions following the World Health Organization’s warning in 2009 of a possible H1N1 flu pandemic are an example of proactive safety management. After the warning was issued, European and other governments began to stockpile considerable amounts of drugs and vaccines to ensure that the necessary resources were in place. Although it later turned out to have been a false alarm, it illustrates the essential features of proactive safety management. It is obviously a problem for proactive safety management that the future is uncertain and that an expected situation may fail to happen. In that case, preparations will have been made in vain, and time and resources may have been wasted. It is also a problem that predictions may be imprecise or incorrect, so that the wrong preparations are made. Proactive safety management thus requires taking a risk, not least an economic one. But the alternative of not being ready when something serious happens will indubitably be even more expensive in both the short and the long run.

5 Conclusion

While day-to-day activities at the sharp end never are reactive only, the pressure in most work situations is to be efficient rather than thorough. This reduces the possibilities to be proactive^[5]. Proactive safety management does require that some effort is spent up front to think about what could possibly happen, to prepare appropriate responses, to allocate resources, and make contingency plans.

In practice, it is easier to be proactive for large-scale

events than for small-scale ones, because they develop relatively slowly – even though they may begin abruptly. Large scale events are regular rather than irregular, and there are often clear indicators for when a response is needed. The appropriate responses are furthermore known, so that preparations can be made ahead of time.

It is more difficult to be proactive for the myriad of small-scale events that constitute everyday work situations. Here, things may develop rapidly and unexpectedly, there are few leading indicators, and resources are often stretched to the limit. There will both be fewer resources to allocate, and less time to deploy them. The pace of work leaves little opportunity to reflect on what is happening and to act strategically. Indeed, work pressures and external demands often lead to opportunistic solutions that force the system into a reactive mode. To get out of this – to switch from a reactive to a proactive mode – requires a deliberate effort. While this may not seem to be affordable in the short term, it is unquestionably a wise investment in the long term.

Here are some practical suggestions for how to begin that process:

- Look at what goes right, as well as what goes wrong. Learn from what succeeds as well as from what fails. Indeed, do not wait for something bad to happen but try to understand what actually took place in situations where nothing out of the ordinary seemed to happen. Things do not go well because people simply follow the procedures. Things go well because people make sensible adjustments according to the demands of the situation. Find out what these adjustments are and try to learn from them!
- When something has gone wrong, look for everyday performance variability rather than for specific causes. Whenever something is done, it is a safe bet that it has been tried before. People are quick to find out which performance adjustments work and soon come to rely on them – precisely because they work. Blaming people for doing what they usually do is therefore counterproductive. Instead one should try to find the performance adjustments people usually make as well as the reasons for them. Things go wrong for the same reasons that they go right, but it is far easier and less incriminating to study how

things go right.

- Look at what happens regularly and focus on events based on how often they happen (frequency) rather than how serious they are (severity). It is much easier to be proactive for that which happens frequently than for that which happens rarely. A small improvement of everyday performance may count more than a large improvement of exceptional performance.
- Allow time to reflect, to learn, and to communicate. If all the time is used trying to make ends meet, there will no time to consolidate experiences or replenish resources – including how the situation is understood. It must be legitimate within the organisational culture to allocate resources – especially time – to reflect, to share experiences, and to learn. If that is not the case, then how can anything ever improve?
- Remain sensible to the possibility of failure – and be mindful. Try to think of – or even make a list of – undesirable situations and imagine how they may occur. Then think of ways in which they can either be prevented from happening, or be recognised and responded to as they are happening. This is the essence of proactive safety management.

5.1 The way ahead

The main reason for juxtaposing Safety-I and Safety-II is to draw attention to the consequences of basing safety management on one or the other. The basic differences are summarised in Table 1 below.

Table 1 Basic difference between Safety-I and Safety-II

	Safety-I	Safety-II
Definition of safety	That as few things as possible go wrong	That as many things as possible go right
Safety management principle	Reactive, respond when something happens	Proactive, try to anticipate developments and events
Explanations of accidents	Accidents are caused by failures and malfunctions	Things basically happen in the same way, regardless of the outcome.
View of the human factor	Liability	Resource

What people do in everyday work situations is usually a mixture of Safety-I and Safety-II. The precise balance depends on many things, such as the nature of the work, the experience of the people, the organisational climate, management and customer pressures, *etc.* Everybody knows that prevention is better than cure, but the conditions may not always be conducive to that.

It is a different matter when it comes to the levels of management and regulatory activities. Here it is clear that the Safety-I view dominates, for reasons that have been explained in the beginning of this note. (The imbalance may be due to an efficiency-thoroughness trade-off as well: it is much simpler to count the few events that fail than the many that do not. And it is also – wrongly – assumed to be easier to explain the former rather than the latter.)

Since the socio-technical systems on which our existence depends continue to become more and more complicated, it seems clear that staying with a Safety-I approach will be inadequate in the long; that may by now even be the case also in the short run. Taking a Safety-II approach should therefore not be a difficult choice to make. Yet the way ahead does not lie in a wholesale replacement of Safety-I by Safety-II, but rather in a combination of the two ways of thinking. It is still the case that the majority of adverse events are relatively simple – or can be treated as relatively simple without serious consequences – and that they therefore can be dealt with in the way we have become accustomed to. But there is a growing number of cases where this approach will not work. For these, it is necessary to adopt a Safety-II perspective – which essentially means adopting a resilience engineering perspective. Safety-II is first and foremost a different way of looking at safety, hence also a different way of applying many of the familiar methods and techniques. In addition to that it will also require methods on its own, to look at things that go right, to analyse how things work, and to manage performance variability rather than just constraining it. We cannot make things go right simply by preventing them from going wrong. We can only make things go right by understanding the nature of everyday performance and by perceiving those things

which we otherwise do not see.

Appendix: Footnote of the term in the text

*1: The English word *safe* comes from the French word *sauf*, which means both ‘without’ and ‘unharmful.’ The origin of the word is the Latin *salvus*, meaning uninjured, healthy, and safe.

*2: Human Factors Engineering was by then more than 30 years old, but had on the whole focused more on productivity issues than safety issues.

References

- [1] HALE, A. R. and HOVDEN, J.: Management and culture: the third age of safety. A review of approaches to organizational aspects of safety, health and environment. In A. M. Feyer & A. Williamson (Eds), Occupational Injury. Risk Prevention and Intervention. London, Taylor & Francis, 1998.
- [2] HARRIS, J. D.: Habitatory response decrement in the intact organism. Psychological Bulletin. 40, 1943: 385–422.
- [3] THOMPSON, R. F. and SPENCER, W. A.: Habituation: A model phenomenon for the study of neuronal substrates of behavior. Psychological Review, 73(1), 1966: 16-43.
- [4] JAMES, W.: The principles of psychology. London, Macmillan and Co, 1890.
- [5] HOLLNAGEL, E.: The ETTO principle. Efficiency-Thoroughness Trade-Off or why things that go right sometimes go wrong. Farnham, UK, Ashgate, 2011.
- [6] AMALBERTI, R.: Optimum system safety and optimum system resilience: Agonistic or antagonistic concepts? In E. Hollnagel, D. D. Woods & N. Leveson (Eds.), *Resilience engineering, Concepts and precepts*. Aldershot, UK, Ashgate, 2006.
- [7] ORGANISME D’ENQUÊTE POUR LES ACCIDENTS ET INCIDENTS FERROVIAIRES: Rapport d’enquete de securité: La collision ferroviaire survenue le 15 fevrier 2010 a Buizingen. Service public fédéral Mobilité et Transports: Bruxelles, Belgium, 2012.
- [8] REASON, J. T.: The human contribution. Farnham, UK, Ashgate, 2008.
- [9] MELCHERS, R. E.: On the ALARP approach to risk management. Reliability Engineering & System Safety, 71(2), 2001: 201–208.
- [10] HOLLNAGEL, E. and WOODS, D. D.: Joint cognitive systems: Foundations of cognitive systems engineering. Boca Raton, FL, CRC Press, 2005.
- [11] WEARS, R. L. and PERRY, S. J.: Free fall - a case study of resilience, its degradation, and recovery, in an emergency department. Paper presented at the 2nd International Symposium on Resilience Engineering, Juan-les-Pins, France, 2006.
- [12] ANTONIO, G. E., GRIFFITH, J. F. and AHUJA, A. T.: Aftermath of SARS. In A. T. Ahuja & C. G. C. Ooi (Eds.): *Imaging in SARS*. Cambridge University Press, 2004: 159-164.
- [13] HOLLNAGEL, E. (Ed.): *Safer complex industrial environments*. Boca Raton, FL, CRC Press, 2010