Available online at www.sciencedirect.com
ScienceDirect

Safety Science 46 (2008) 221–229

SAFETY SCIENCE

www.elsevier.com/locate/ssci

Risk + barriers = safety?

Erik Hollnagel

École des Mines de Paris – Pôle Cindyniques, Sophia Antipolis, France

Received 13 November 2006; received in revised form 10 June 2007; accepted 11 June 2007

Abstract

According to a common safety model, safety can be brought about either by eliminating hazards, by preventing initiating events, and/or by protecting against outcomes. The two primary types of responses, prevention and protection, both involve the use of barriers in one way or another. The paper discusses the characteristics of different barrier systems (physical, functional, symbolic, and incorporeal) and their relative advantages and disadvantages. It is argued that while barriers are necessary, they basically represent a reactive approach which is insufficient by itself to guarantee safety.

Keywords: Risk; Safety; Barriers; Substitution principle; Safety management

1. Introduction

Risk and safety are linked both conceptually and pragmatically. The conceptual link can be seen by comparing definitions of the two concepts. Risk, for instance, is normally defined as the likelihood that something unwanted can happen. Safety is likewise defined as the absence of unwanted events, which essentially means as the absence of risk. The pragmatic link mirrors this reciprocity as seen from the fact that safety – or rather, the lack thereof – usually is measured by the number of specified unwanted events, such as accidents and incidents. A higher level of safety is equivalent to a lower occurrence of such events and therefore, to a lower level of risk.

It is a near unavoidable consequence of these definitions that the best way to ensure a state of safety is either to prevent something unwanted from happening or to protect against its consequences, as illustrated by Fig. 1. Since, in practice, it is impossible completely to prevent unwanted events, i.e., completely to eliminate risks, the two approaches are best used together.

1.1. Safety through risk elimination

Although prevention in many ways is better than protection, it is a fact of life that perfect prevention is impossible. This realisation has been made famous by the observation that there always is something that

E-mail address: erik.hollnagel@cindy.ensmp.fr

0925-7535/$ - see front matter © 2007 Elsevier Ltd. All rights reserved.
doi:10.1016/j.ssci.2007.06.028

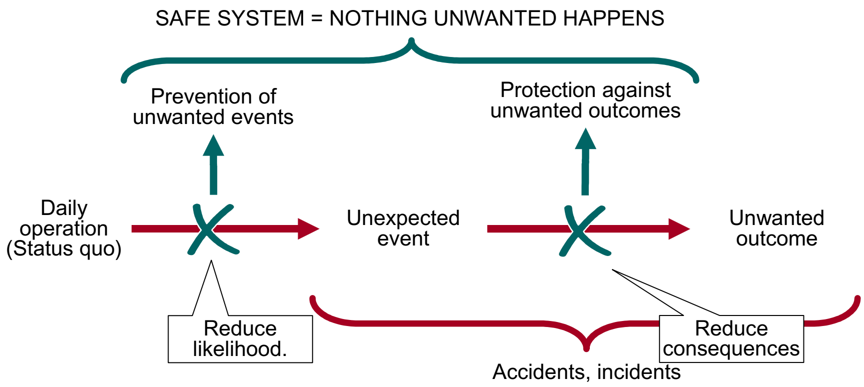
SAFE SYSTEM = NOTHING UNWANTED HAPPENS
Prevention of unwanted events
Protection against unwanted outcomes
Daily operation (Status quo)
Unexpected event
Unwanted outcome
Reduce likelihood.
Reduce consequences
Accidents, incidents

Fig. 1. Safety through prevention and protection.

can go wrong. Although the anonymous creator of this truism never will be known, it is certain to have been uttered millennia before either Josiah Spode (1733–1797) or the hapless Major Edward A. Murphy Jr. (In passing there is also Ambrose Bierce's definition of an accident as "(a)n inevitable occurrence due to the action of immutable natural laws".) A more sophisticated version of that is Perrow's (1984) thesis that systems by the 1980s had become so complex that accidents should be considered normal events.

In order to ensure safety by preventing something from happening, i.e., through the elimination of risks, it is first of all necessary that the risks are known or can be made known. To do so is the purpose of risk assessment, and there are a considerable number of well-established methods available for that (Aven, 2003; Leveson, 1995). These methods usually combine a *representation* of how events may develop, where event and fault trees are characteristic examples, with ways of *estimating* and/or *calculating* the probability that a specific event or combination of events obtain. To avoid the pitfalls of relying on routine and established norms, risk assessment requires a certain level of imagination, as argued by Adamski and Westrum (2003). This is, however, not a topic that will be addressed here.

The pursuit of safety through the elimination of risks also required that the specific risk source actually can be removed from the system without impeding or changing the system's functioning. In some cases, this condition is obviously violated as when the elimination of a risk means the loss of a primary function. Thus, the risk of an airplane falling down can only be fully eliminated by not taking to the air, but that is clearly not a viable option, at least in commercial aviation. (Note, however, that it may be acceptable in other cases, e.g., the grounding of the space shuttle after the Columbia accident). In other cases where the risk is eliminated by substituting one function for another, the condition is apparently met. Unfortunately this is not actually so, the reason for that being that the substitution principle is invalid.

1.2. The risk of the substitution principle

Elimination by means of substitution is an often used solution, the most conspicuous case being when human performance is replaced by technology, specifically by automation. The rationale for this is that automation is highly reliable because it is the result of a formal design process and because it is based on components with known failure rates. Humans, on the other hand, are generally seen as fallible and unreliable, as 'proved' by countless examples of 'human error'. The fallacy of this argument should by now be so obvious that it hardly needs to be belaboured.

The substitution principle expresses the common assumption that artefacts are neutral in their effects and that their introduction into a system therefore only has intended and no unintended consequences. The basis for this principle is the concept of interchangeability, which of course has proved its value as the basis for large scale industrialisation. Thus, if there are a number of identical parts, such as light bulbs or pumps, it is possible to replace one by another without unwanted side-effects. In general, however, substitutability only works when parts are not interacting and when there is no appreciable tear and wear. If parts are interacting, they constitute a system with dependencies, which almost by definition invalidates the substitution assumption.

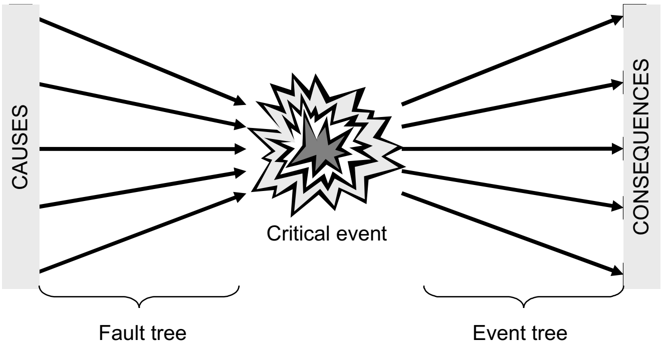

Fig. 2. The bowtie model.

If the substitution principle is dubious in the case of technological components, it is even more suspect in the case of a substitution of functionality of which the extreme is replacing humans by automation. Carroll and Campbell (1988), for instance, noted that "(n)ew tools alter the tasks for which they were designed, indeed alter the situations in which the tasks occur and even the conditions that cause people to want to engage in the tasks" (p. 4; see also Sarter et al., 1997). In relation to risk reduction, this means that a substitution of functions changes the basis for risk assessment, often in a fundamental way. It is consequently not warranted to claim that the substitution leads to a reduction in risk unless both short- and long-term consequences of the change have been fully taken into account.

1.3. Safety through prevention and protection

Whereas, Fig. 1 illustrates the high-level principles for accident avoidance, hence safety, the aetiology of accidents, hence the nature of risk, is more intricate. Fig. 2 shows the so-called bowtie model, which describes how a critical event may have several precursors, as well as several consequences (Delvosalle et al., 2005). Going through Fig. 2 from left to right suggests at least three different ways of achieving safety.

The first option is to prevent the critical event from taking place. According to the logic of Fig. 2, and therefore, also the logic of the bowtie model, this can be done by hindering preconditions or initiating factors from having an effect that changes the critical event from a possibility to a reality. Note that this does not necessarily require the elimination of these factors (which would be applying the basic principle recursively), but rather that the functions in question are rendered ineffectual.

The second option is to eliminate the critical event altogether, either directly or by substitution, as discussed above. The third option is to protect against the consequences of the critical event if or when it happens, all the precautions notwithstanding. This can be done by reducing or weakening the consequences or by changing their direction either in a real or in a metaphorical sense. Note that, whereas, the first option, prevention, tries to maintain the functioning of the system and to keep it going, the third option, protection, does not need to do that. Indeed, protection may require that the system is shut down when the critical event occurs, as in the case of nuclear power plants, or that the normal functioning is reduced until the situation again has returned to normal. This is acceptable because the goal in such cases is not the continued functioning but the safety of the larger system, such as the general population or the environment.

2. Responding to accidents

It is a universally accepted adage that prevention is better than cure. In what we may broadly call the industries, ranging from production and construction to transportation and communication, safety related resources are therefore spent mostly on prevention. The main focus is on actual events, on the accidents that have happened and the aim is to avoid a repetition or recurrence of them. Safety efforts ought, however, also consider the potential events, hence look at what might happen in the future. Many industry people, unfortunately, often find it hard in practice to weigh a certain expense (i.e., the cost of prevention) against an uncer-

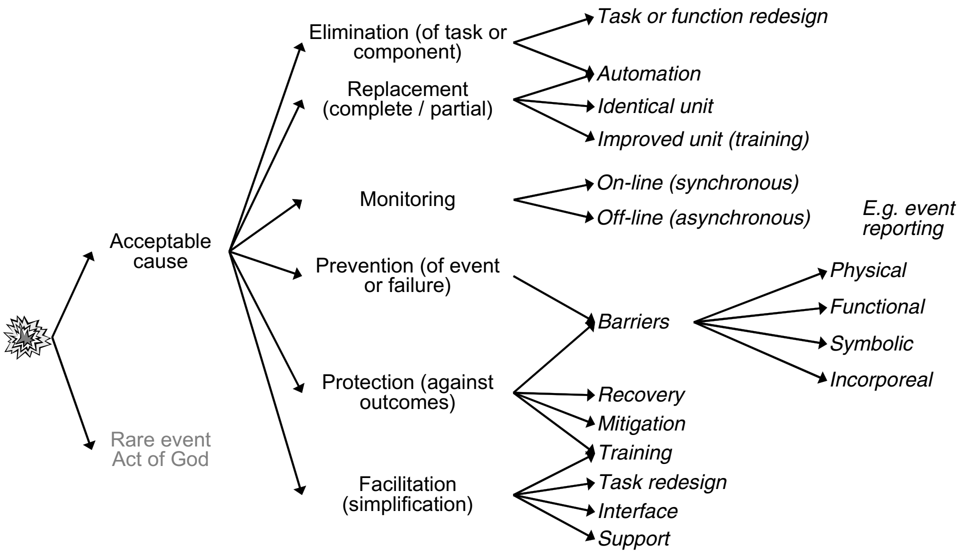

Fig. 3. Characterisation of typical responses to accidents.

tain gain (i.e., an accident avoided) and therefore, prefer to concentrate on what has happened rather than on what possibly may happen.

Since safety efforts usually are motivated by accidents and incidents that have taken place, it is instructive to consider the range of responses to accidents. These can be categorised as shown by Fig. 3, the variety of approaches to safety discussed above obviously being among them.

Table 1
Possible reactive strategies for accident prevention

Main strategy	Type	Example
Elimination	Cancellation	Withdrawing a product from the market
	Restructuring	Making a function unnecessary through redesign
Replacement (complete or partial)	Identical unit or component	Spare parts or components; backups
	Improved unit or component	New models; new software releases; automation
Monitoring	Early warnings	Performance indicators; alerts and alarms
Prevention	Functional barrier system	Alarms
		Interlocks
		Interface
	Physical barrier system	Buildings, fences
	Symbolic barrier system	Rules, tasks
		Procedures
	Incorporeal barrier system	Safety culture
Protection	Physical barrier system	Wall
	Functional barrier system	Airbag
	Recovery	System design
		Operational support
		Fault tolerance
	Mitigation	Feedback
		Detection
		Undoing
Facilitation	Task redesign; work design	Improved task ‘logic’; collaborative work
	Interface design	Consistency; usability; functional grouping
	Support	Attention, memory

When considering how to respond to an accident, it is first of all necessary to make clear whether what happened was something that can be expected to recur or whether it was a 'freak' event. In the latter case, the response may be a decision to do nothing, either because no reasonable explanation can be constructed using the current scientific vocabulary and set of beliefs (meaning that it is considered as an 'act of god') or because calculations show that the event is so rare that it is not worthwhile to do anything about it. This is essentially the reasoning behind PSA/PRA, where the main purpose is to determine the probabilities of future, risky events and from this basis make the changes necessary to ensure that the probabilities are too low to be of any concern.

Normally it is, however, necessary to do something and the responses can be in terms of elimination, replacement, monitoring, prevention, protection, or facilitation (cf. Fig. 3; see also Table 1 and the description in Hollnagel, 2004). *Elimination* is the complete removal of the offending system or component, as discussed above. This is a relatively rare occurrence, because it either means the loss of a major investment – to say nothing of prestige – or a significant cost to restructure or reconfigure the rest of the system. *Replacement* can either be by an identical component or by one that has been an improved in some way. In the case of humans, the 'identical component' can be another person with the same nominal qualifications but without a history of failure while an 'improved component' can be a person who have been to a training course. *Monitoring* does not as such improve safety, but it increases the likelihood that accidents may be detected at an early stage, hence that the consequences will be less severe. *Prevention* is the blocking or hindering of preconditions or initiating factors from triggering or contributing to an accident. As will be discussed below, prevention is always accomplished by using the four possible barrier systems, either individually or in combination. *Protection* is the blocking or hindering of consequences and involves the use of either physical or functional barrier systems. *Facilitation*, finally, is the solution to change or modify the system and/or the working conditions so that it becomes more difficult to do something incorrectly. Facilitation often relies on interface design, procedures and training.

The two primary types of responses, prevention and protection, both involve the use of barriers in one way or another. (In a more metaphorical sense, this is also the case for monitoring and facilitation.) The term *barrier* is often used haphazardly, and in many cases, various forms of barriers are invented as the need arises. Examples are social barriers, organisational barriers, hardware barriers, cultural barriers, behavioural barriers, human barriers, etc. As already Svenson (1991) pointed out, it is not only useful but outright necessary to use a more precise terminology and in particular to make a distinction between barrier systems and barrier functions.

3. Barrier systems and barrier functions

A starting point for a terminology is a distinction between what barriers *do*, i.e., their purpose or function, and what barriers *are*, i.e., the way(s) in which they achieve their purpose. The first, called the barrier functions, describes the modes by which it is possible generically to prevent, or protect against, the uncontrolled transportation of mass, energy, or information. The second, called the barrier systems, describes the means by which the barrier functions are carried out. In order to characterise the possible barrier systems, it turns out that the following four types are sufficient (see also Table 2).

- *Physical* or *material barrier systems*, which concretely prevent an event from taking place or mitigate the effects of an unexpected event by blocking the transportation of mass, energy or information from one place to another. Examples of physical barrier systems are buildings, walls, fences, railings, bars, cages, gates, containers, fire curtains, etc. An important characteristic of physical barrier systems is that they do not have to be perceived or interpreted by someone (or something) in order to work. They can therefore be used against energy and material, as well as against people.
- *Functional barrier systems* create one or more pre-conditions that have to be met before an action can be carried out, for instance by establishing an interlock, either logical or temporal (cf. Leveson, 1995). Some functional barrier systems require a user to change from one state to another; others are autonomous and can change depending on external conditions. A functional barrier system may not always be visible or discernible to a human user, although its presence usually is indicated in some way.

Table 2
Barrier functions for the four barrier systems

Barrier system	Barrier function	Example
Physical	Contain or protect. Prevent transporting something from the present location (release) or into another (intrusion)	Walls, doors, buildings, restricted physical access, railings, fences, filters, containers, tanks, valves, rectifiers, etc.
	Restrain or prevent movement or transportation of mass or energy	Safety belts, harnesses, fences, cages, spatial distance (gulfs, gaps), etc.
	Keep together. Cohesion, resistance	Components that do not break easily (safety glass).
	Separate, protect, block	Crumble zones, scrubbers, filters, etc.
Functional	Prevent movement or action (mechanical, hard)	Locks, equipment alignment, physical interlocking, equipment match, etc.
	Prevent movement or action (logical, soft)	Passwords, entry codes, action sequences, pre-conditions, physiological matching, etc.
	Hinder or impede actions (spatio-temporal)	Distance, persistence, delays, synchronisation, etc.
	Dampen, attenuate	Active noise reduction, active suspension
	Dissipate energy, quench, extinguish	Air bags, sprinklers, etc.
Symbolic	Counter, prevent or thwart actions (visual, tactile interface design)	Coding of functions, demarcations, labels & warnings (static), etc.
	Regulate actions	Instructions, procedures, dialogues, etc.
	Indicate system status or condition (signs, signals and symbols)	Signs (e.g., traffic signs), signals (visual, auditory), warnings, alarms, etc.
	Permission or authorisation (or the lack thereof)	Work permit, work order
	Communication, interpersonal dependency	Clearance, approval, (on-line or off-line), in the sense that the lack of clearance, etc., is a barrier
Incorporeal	Comply, conform to	Self-restraint, ethical norms, morals, social or group pressure
	Prescribing: rules, laws, guidelines, prohibitions	Rules, restrictions, laws (all either conditional or unconditional), etc.

- *Symbolic barrier systems* work indirectly through their ‘meaning’, hence require an act of interpretation by someone. Symbolic barrier systems are ubiquitous in a modern society and we are permanently surrounded by a variety of visual and auditory signs and signals, warnings (by text or by symbol), alarms, etc.
- *Incorporeal barrier systems*, which are not physically present in the situations where they are applied but depend on the knowledge of the user to achieve their purpose. In industrial contexts, incorporeal barrier systems are largely synonymous with the so-called organisational barriers, i.e., rules for actions that are imposed by the organisation, rather than being physically, functionally or symbolically present in the system.

3.1. Organisational barriers

Discussions about safety frequently invoke the concept of an ‘organisational barrier’. Leaving aside the fact that an organisation well may be a barrier in a metaphorical sense, the reason for the popularity of the concept is probably that barriers often are initiated and implemented by an organisation. Yet the actual barrier is in the form of rules or procedures that are carried out by people rather than the organisation as such. Consider, for instance, the case of a work permit. Not having a work permit is supposed to prevent work from taking place. In many cases, work nevertheless starts before permission is given, either because it is expected to be granted, or because it is hoped that a *fait accompli* in the end will lead to the work permit being issued. This example shows that a work permit is not a functional barrier system since there is nothing in the executing system itself that prevents the actions from being carried out, short of the ethics or morals (or fear of punishment) of the people involved. In the classification used here, a work permit rather represents a symbolic barrier system, since it is a token and since the efficiency depends on interpretation. It is therefore not sensible to classify a work as an organisational barrier, even though it comes from the organisation.

3.2. Composite barrier systems

In order to be effective, barriers often rely on a combination of barrier systems. For instance, general speed limits (incorporeal) given by the traffic laws can be supplemented by speed bumps (physical), road signs (symbolic) and at times enforced by traffic police (acting as a symbolic barrier system if they are seen and as a functional barrier system if they are unseen). Symbolic and incorporeal barrier systems must usually be amplified by physical and functional barrier systems to work. Physical and functional barrier systems may likewise be complemented by symbolic barrier systems to encourage their use.

In the case of physical and functional barrier systems, the respective barrier function is provided by the system itself. In the case of symbolic and incorporeal barrier systems, the barrier systems cannot themselves provide the barrier function but require an action by someone (or possibly something). The action thus implements the barrier function, which in turn is provided by the barrier system. This is not just a play with words since knowing which barrier functions an action serves and whether the action came from an instruction (a symbolic barrier system) or a sense of duty (an incorporeal barrier system) can have significant consequences.

3.3. Barrier functions

Whereas, the barrier systems seem to be of just four types – allowing, of course, for combinations among them – barrier functions are less easy to put in order, the attempts of philosophers notwithstanding (e.g., von Wright, 1971). One possibility is to distinguish whether the barrier function is active or passive, i.e., whether it *does* something, such as sprinkler (functional barrier system) that extinguishes a fire, or whether it simply *is*, such as a wall (physical barrier system) which blocks the transportation of matter and energy. (Note, by the way, that neither a symbolic nor an incorporeal barrier system would be able to protect against a fire once it had started.) But as Table 2 shows, the same barrier function can be realised in different ways, i.e., with different barrier systems as the basis.

The classification of barriers is unfortunately not always simple as the examples given above. A wall is, of course, a physical barrier system and a law is an incorporeal barrier system. But how does one characterise a procedure? A procedure is an instruction for how to do something and is therefore an example of facilitation rather than prevention. Procedures do, however, often include both cautions and conditional actions (if–then rules). The procedure also works by virtue of its contents or meaning rather than by virtue of its physical characteristics. For that reason, it is therefore warranted to classify a procedure as a symbolic barrier system.

Although it is analytically correct to maintain a clear distinction between barrier systems and barrier functions, it may be cumbersome in practice since it leads to clumsy expressions. In cases where it does not create problems it is therefore practical just to talk about a barrier, as a short-hand reference to a barrier function implemented by a barrier system.

3.4. Conditions for effective barriers

Most accidents are due to a combination of an unexpected event and a dysfunctional or missing barrier, rather than to a single initiating action (cf., Reason, 1990). It is therefore essential that barriers can achieve their purpose, both in theory and in practice. In order to determine whether this is the case, it is useful to look at the set of conditions that must be fulfilled. These are summarised in Table 3 below, which also provides an initial evaluation of how different barrier systems fare with respect to the conditions.

Incorporeal barrier systems generally score low on all evaluation categories because they completely depend on the users' willingness to abide by them. They may nevertheless be attractive to managers both because the resource needs are low and because the delay in implementation short. Unless the population of users have unusually high moral standards, incorporeal barrier systems are not advisable, except as a temporary remedy. There may possibly be a relation between incorporeal barrier systems and safety culture, but this discussion must be saved for another time.

Table 3
Evaluation of barrier system quality

Quality	Short definition	Physical	Functional	Symbolic	Incorporeal
Efficiency	How well the barrier meets its intended purpose	High	High	Medium	Low
Resource needs (cost)	What is needed to design, develop and maintain a barrier	Medium–high	Low–medium	Low–medium	Low
Robustness (reliability)	How well a barrier can withstand the variability of the environment.	Medium–high	Medium–high	Low–medium	Low
Implementation delay	The time from conception to implementation of a barrier	Long	Medium–long	Medium	Short
Applicable to safety critical tasks	Self-explanatory	Low	Medium	Low (uncertain interpretation)	Low
Availability	Whether a barrier can fulfil its purpose when needed. This is critical for barriers designed for rare conditions	High	Low–high	High	Uncertain
Evaluation	How easy it is to determine whether a barrier works as expected, both during design and actual use	Easy	Difficult	Difficult	Difficult
Independence on humans (during operation)	The extent to which a barrier does not depend on humans to achieve its purpose	High – and in principle completely independent	High	Low	Low

Symbolic barrier systems are relatively inexpensive and can be put in place rather quickly. On the other hand, their efficiency is low, since people can choose simply to ignore them (the dire warnings on cigarette packs are a case in point). They are therefore ill-suited for safety-critical tasks, at least as the only barrier system. Another drawback is that people quickly become used to them, hence stop paying attention unless a violation is associated with a significant subjective risk or backed by strong negative reinforcement.

Functional barrier systems are on the whole very efficient, but often require complex preparations and are therefore, both costly and lengthy to implement. This is particularly so if they are introduced well after the system has been built, e.g., in response to a newly discovered risk. They can be very reliable, although this in the case of hardware systems may require extensive maintenance (software is obviously much easier in that respect, since it does not degrade over time). Functional barrier systems can also be difficult to evaluate, either because the test and evaluation itself may affect the barrier or because the preconditions and dependencies may be quite complex.

Physical barrier systems finally score high on most criteria, i.e., they are efficient, robust, independent of humans, and generally easy to verify. Their disadvantage is that they can be very costly and time-consuming to establish, and also that they may require considerable maintenance. They are therefore most useful for risks that have been identified as part of the system design process, or for risks that warrant a significant system redesign.

Effective barriers normally rely on a mixture of barrier systems, with the possible exception of some physical barrier systems (walls, gates, moats, etc.). Several barriers may also be used together to increase the robustness, not least for safety-critical applications. Typical examples are the use of defense-in-depth, such as in nuclear power plants and in modern automobiles. Having more than one level of barriers obviously reduce the probability that an accident will happen, although it can never completely eliminate it. As an overriding concern, it is important to be able to assess the vulnerability of the barriers that are put in place as well as how they may fail – singly or in combination. Accident analysis does that by determining which barriers failed and why they failed. System design and risk assessment do that by identifying how barriers may potentially fail, i.e., what the weaknesses of the system are. The descriptions of vulnerabilities may be further developed into an assessment of the reliability of the different barrier systems. Such an assessment may be a valuable input to system design.

4. Conclusions

Although it is not a necessary trait of barriers, most of them are used to prevent something from occurring again. In other words, barriers are used as a reaction or a response. Safety can however, not be guaranteed only by reacting. It is equally important to look ahead, to identify potential new risks, and then to devise barriers against them.

Westrum (2006) has proposed a distinction between regular threats, irregular threats, and unexampled events. *Regular threats* are events that occur so often that the system can develop a standard response. An example is medication errors that only implicate a single patient, and which potentially can be brought under control. *Irregular threats* are one-off events where their sheer number makes it practically impossible to provide a standard response. Even though they are imaginable, they are usually unexpected. Finally, *unexampled events* are those that are virtually impossible to imagine and which exceed the responders' collective experience. Since irregular threats and unexampled events are infrequent and unusual, they cannot be treated in the conventional way, i.e., by designing barriers to avoid them. Indeed, they cannot easily be described by the linear types of accident models that are *de rigueur* in safety management. Their distinguishing feature seems to be that they emerge out of a situation. The proper way to deal with them is therefore to address the situations or conditions where they can occur. This is not primarily achieved by the use of barriers, but rather by alternative techniques such as performance variability management or resilience engineering (Hollnagel et al., 2006).

Barriers are an effective means against known risks, a way to prevent unwanted events from taking place and to protect against their consequences. Yet barrier design must not become entirely reactive, since in that case safety will become a game of constant fire fighting or catching up. Safety cannot genuinely be improved only by looking to the past and taking precautions against the accidents that have happened. Safety must also look to the future. It must be proactive, although that requires taking a risk – the unwanted outcome being that nothing untoward happens and that the investment therefore is not matched by tangible results.

References

Adamski, A.J., Westrum, R., 2003. Requisite imagination. The fine art of anticipating what might go wrong. In: Hollnagel, E. (Ed.), Cognitive Task Design. Lawrence Erlbaum Associates, Mahwah, NJ.

Aven, T., 2003. Foundations of Risk Analysis. A Knowledge and Decision-oriented Approach. John Wiley and Sons, Ltd., Chichester.

Carroll, J.M., Campbell, R.L., 1988. Artifacts as Psychological Theories: The Case of Human–Computer Interaction. User Interface Institute, IBM T.J. Watson Research Centre, Yorktown Heights, NY.

Delvosalle, C., Fiévez, C., Pipart, A., Casal Fabrega, J., Planas, E., Christou, M., Mushtaq, F., 2005. Identification of reference accident scenarios in SEVESO establishments. Reliability Engineering and System Safety 90, 238–246.

Hollnagel, E., 2004. Barriers and Accident Prevention. Ashgate Publishing Limited, Aldershot, UK.

Hollnagel, E., Woods, D.D., Leveson, N.G. (Eds.), 2006. Resilience Engineering: Concepts and Precepts. Ashgate Publishing Limited, Aldershot, UK.

Leveson, N.G., 1995. Safeware. System Safety and Computers. Addison-Wesley, Reading, MA.

Perrow, C., 1984. Normal Accidents: Living with High Risk Technologies. Basic Books, Inc., New York.

Reason, J.T., 1990. Human Error. Cambridge University Press, Cambridge, UK.

Sarter, N.B., Woods, D.D., Billings, C.E., 1997. Automation surprises. In: Salvendy, G. (Ed.), Handbook of Human Factors and Ergonomics, second ed. Wiley, NY.

Svenson, O., 1991. The accident evolution and barrier function (AEB) model applied to incident analysis in the processing industries. Risk Analysis 11 (3), 499–507.

von Wright, G.H., 1971. Explanation and Understanding. Routledge and Kegan Paul, London.

Westrum, R., 2006. A typology of resilience situations. In: Hollnagel, E., Woods, D.D., Leveson, N.G. (Eds.), Resilience Engineering: Concepts and Precepts. Ashgate Publishing Limited, Aldershot, UK.