Contents lists available at ScienceDirect

Process Safety and Environmental Protection

journal homepage: www.elsevier.com/locate/psep

Human factors in barrier management: Hard truths and challenges

Ronald W. McLeod

Ron McLeod Ltd., 16 Waterside Avenue, Glasgow, G77 6TJ Scotland, UK

ARTICLE INFO

Article history:
Received 20 October 2016
Received in revised form 8 January 2017
Accepted 12 January 2017
Available online 20 January 2017

Keywords:
Human factors
Barrier management
Bowtie analysis
Human and organisational factors
Human factors engineering
Cognitive bias

ABSTRACT

Whether acting as controls in their own right, being relied on to ensure physical, hardware or electronic controls are in place and functional when needed, or as a threat to safe operations, human performance is central to the development, implementation and sustainable operation of barrier management systems. Many organisations however struggle to know how to ensure: (a) that the human performance they rely on can reasonably be expected to happen when and where it is needed; and, (b) that the controls they intend to have in place are as robust as they reasonably can be to the loss of human reliability. Drawing on real-world incidents, this paper examines some of the expectations that are widely held about human performance in barrier management systems. Those expectations are considered in the light of the reality behind how people think, behave and perform in real world tasks: what are referred to as "hard truths" of human performance. Drawing largely on the technique of Bowtie analysis, weaknesses in the way human factors are treated in current approaches to barrier management are reviewed, and improvements suggested. The paper illustrates how most human and organisational controls should be treated as safeguards rather than full barriers: they are critical to safety management, but are rarely able to meet the criteria necessary to be treated as full barriers. The equivalence of a Human Performance Standard to an Equipment Performance Standard is illustrated with a practical example.

1. Introduction

The importance of having in place a number of controls (also variously referred to as "layers of defences", "protective measures", "barriers", and "protection layers") to protect against the risk of major losses is now virtually ubiquitous across industries with the potential for major accidents. International standards such as IEC 61508 (IEC, 2003) and 61511 (IEC, 2010), as well as a variety of sources of guidance on best practice (PSLG, 2009; CCPS, 2015, 2017; CIEHF, 2016) set out recommended approaches to developing, analysing and validating layers-of-defences strategies. And there is a growing literature on experience and lessons learned with different techniques (see for example, Sklet, 2006; Chambers et al., 2009; Lewis and Smith, 2010; Hamilton and Turner, 2014).

Human performance continues to be relied on for controls to do what is expected: whether the decisions and actions of people act as controls in their own right, or whether they are relied on to ensure physical, hardware or electronic controls are in place and functional when needed. On the other hand, loss of human reliability – "human error" – is widely regarded as one of the principal threats that need to be guarded against through the use of barrier models. One of the main findings of a review of the application of Layers-of-Protection-Analysis (LOPA) to the risk of overspill at fuel storage tanks in the UK, was that: "*Human factors appear to dominate a number of initiating event (IE) frequencies and conditional modifier (CM) error probabilities in all the LOPA studies assessed in this work*". (Chambers et al., 2009, p. 2).

Many organisations however struggle to know how to ensure: (a) that the human performance they rely on can reasonably be expected to happen when and where it is needed; and (b) that the controls they intend to have in place are as robust as they reasonably can be to the loss of human reliability.

This paper has two objectives: first to examine some of the expectations that are widely held about human performance in barrier management systems, and to consider those expectations in the light

E-mail addresses: ron@ronmcleod.com, ronmcleod2@gmail.com
http://dx.doi.org/10.1016/j.psep.2017.01.012

of some of the reality behind how people think, behave and perform in real world tasks; and, second, to review some common weaknesses with the way human factors are treated in many current approaches to barrier management and to suggest some improvements. While much of the content applies to many formalized approaches to barrier management, the paper is based predominantly around the use of Bowtie analysis as a means of identifying and managing barriers (see CCPS, 2017 for an introduction to Bowtie analysis).

2. Expectations of people in barrier systems

The key premise behind this paper is that both safety and productivity would benefit by being clearer about the role of people in operations, and what exactly it is organisations need and expect of people for the safe and reliable operation of their assets. And having achieved that clarity, both ensuring that those expectations are reasonable, and focusing on what needs to be done to enable people to perform to the standards and to the level of consistency and reliability that is needed and expected.

When things go wrong, investigations commonly reach a conclusion that, if only people had followed procedures the incident would not have happened (see for example MMS, 2005). There are a number of implicit expectations underlying such a position: that the organisation has in place all of the procedures it needs; that they are sufficiently specific, accurate, clear and up to date; that they are accessible where and when they are needed; that the people expected to use them have the knowledge, skills and training to know what procedures to use in whatever context; and that they will actually recognise situations where procedures should be followed, identify the correct procedure, and be able to carry them out under the conditions that exist at the time. Such expectations can be far from reality, as many accident investigations have found.

One of the tasks that the crew of the Deepwater Horizon apparently got "wrong" in the final moments before they were overtaken by events on April 20, 2010, was failing to re-set manual controls on a control panel so that fluids returning from the well would be diverted over the side, rather than remaining on-board (CSB, 2016). According to the Chemical Safety Board, had they realised what was happening and taken the necessary action in time, even had the kick occurred, the flammable gas would not have found an ignition source, and the explosion and subsequent losses would not have occurred. Though, as the CSB demonstrated, expectations about the ability of the crew to recognise the situation, make the required decision and take the action to divert the flow, in the available time and with the information they had could not credibly be considered as reasonable for any drill crew.

There is a fundamental conflict between the psychological characteristics of someone who is competent and has expertise, and someone who is expected simply to follow procedures. The expectation that if only people followed procedures, everything would function safely and reliably can be in direct conflict with the requirement that operations will be manned by competent and experienced people: especially when those people are "specialists" or "experts". McLeod (2015, 2016a) emphasises the need to be clear about the real role of people – what is intended and what is expected of them – for controls to operate as intended. And to acknowledge that upsets, incidents and even disasters are often averted because motivated, competent people are able to vary the way they work, including diverging from standard procedures and prescribed work practices. This is the difference between what Hollnagel (2014) describes as "Work-as-Imagined", and "Work-as-Done": "*...everyone at the sharp end knows that it is only possible to work by continually adjusting what they do to the situation*" (Hollnagel, 2014).

Fig. 1 – A walkway on an offshore platform (from McLeod, 2015).

2.1. *Hard-truths of human performance*

A further reason organisations find it so difficult to anticipate, and are so often surprised, by situations where loss of human reliability occurs, is because expectations of how people are likely to behave rarely take into account some of the "hard truths" of human performance (McLeod, 2015). These are truths about how human beings see and interpret the world, think, behave and perform. They are 'hard' because they can be so difficult and inconvenient to design for and to manage. While that is unfortunate, it does not make them any less true: they cannot be dismissed or ignored because they are difficult or inconvenient.

Among the most important of these "hard truths" are:

I Human emotion, thought, performance and attitudes are highly situated—they are strongly influenced by the situation or context as the individuals involved experience and believe it to be at the time;

II The design and layout of work systems, equipment interfaces and the environment influence the ways people behave and interact with technology and the world;

III People will find the easier way of doing things, even if it is more risky, and;

IV People cannot be assumed to be rational.

The third of these is so important that psychologists have even defined a law that governs it: "*The Law of Least Effort...asserts that if there are several ways of achieving the same goal, people will eventually gravitate to the least demanding course of action....laziness is built deep into our nature*" (Kahneman, 2012).

Fig. 1 shows an example of people finding an easier way on an offshore oil platform. The area here provides a walkway from the bottom right of the picture, across the step over plate and on to the top-left. Note the number of footprints on the deck and on the service pipes but not on the step platform. Despite all of the safety training, procedures and efforts to develop a strong safety culture, the workers had found a way to make the route shorter by stepping on the pipes instead of the platform provided. Why? Because it was a few steps shorter. Also, the step platform was approximately 16″ above the deck (the normal riser height is 8″). So a combination of a more

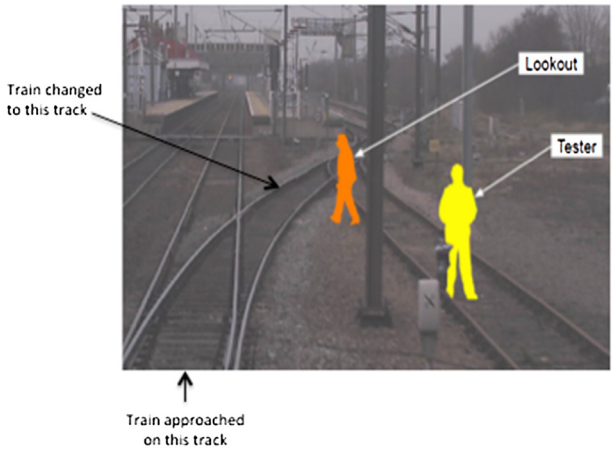

Fig. 2 – Still from video from approaching train (from RAIB, 2015).

direct route, added to a slightly awkward step up, together with apparently no perceived increase in risk from stepping on the pipes, the workers to take the shortcut. Although stepping on pipes certainly carries risk, especially if they are wet, or the workers are in a hurry.

2.1.1. *Styles of thinking, irrationality and cognitive bias*

The fourth hard truth of human performance is that people cannot be assumed to be rational. There is a widespread lack of understanding of the operational significance of what psychologists refer to as two styles of thinking. (See Kahneman (2012) for a comprehensive review of the difference between fast and slow thinking, and of some of the characteristics and biases that can be associated with "slow" thinking). System 1, or "fast" thinking, is intuitive, effortless and rapid. System 2, by contrast, is slow, careful, evidence-based, doubting and rational. System 1 is always "on", while using system 2 needs the application of conscious effort. Most of the time, System 1 is the means by which we are able to think and perform efficiently despite the presence of what would otherwise be an overwhelming array of information, options and uncertainty. Though System 1 has characteristics that make it far from ideal when it comes to achieving high levels of consistently reliable human performance: it is subject to many sources of irrationality or bias, and does not see doubt or ambiguity. It is "*. . .a system for jumping to conclusions*" (Kahneman, 2012).

McLeod (2015, 2016b,c) has explored the implications of these two styles of thinking to operational risk assessment and decision making across all levels of organisations and demonstrated how they can bring insight into why operators may have behaved and taken the actions they did in the events leading up to major accidents.

The implications of System 1 thinking for safety can be illustrated by considering an incident that occurred in the UK in 2014 where a railway lookout walked into the path of an oncoming train (RAIB, 2015). Fig. 2 shows the situation a few moments prior to the incident. The train driver sounded his horn, and the lookout raised an arm in acknowledgement. It was not until about a second before he was struck and killed that the lookout turned and saw the oncoming train. Why would he do that? Why would a trained and experienced operator, who fully understood the risks, and having no desire to cause harm to himself or anyone else, deliberately walk in front of an oncoming train? What was he thinking? Or, more to the point, how was he thinking? While it does not form part of the conclusions of the formal incident investigation, the behaviour of the lookout in this incident has characteristics that are consistent with someone who had assessed the situation and made a decision using System 1 thinking: someone who, when he first saw the train approaching, jumped to a conclusion about the future path of the train into the station. And having made that judgement, had no doubt and did not consider the possibility that the train might – as it in fact did – change track onto the line closer to him. If he had any doubt at all, he would surely have at least looked in the direction of the train when it sounded its final warning. But System 1, Kahneman tells us, does not experience doubt.

2.1.2. *Framing and loss aversion*

It is well known that the cognitive bias known as Framing – presenting the same information in different ways – can evoke very different psychological responses. Framing is particularly important when it is combined with another bias—loss aversion. The science is clear that losses have about twice the psychological impact of equivalent gains: we will work twice as hard to avoid a loss, as we will to achieve the equivalent gain. This is central to Prospect Theory (Tversky and Kahneman, 1974), the work for which Professor Kahneman received his Nobel prize in 2002.

What is the relevance of Prospect Theory to the safety and reliability of industrial operations? Real-world operations do not involve binary choices between personal financial gain and loss. And they are not made by individuals expressing personal choices. But could the psychological preferences described by Prospect Theory apply when front line operations teams make judgements about the real-time risks facing them? Could a psychological preference to work disproportionately harder to avoid what has been framed as a 'loss' than to achieve a 'gain' dominate judgement, thinking, reasoning and decision-making in ways that are not consistent with the risks actually being faced? Could it lead attention, effort and resources to be overly focused on the wrong things?

On April 20, 2010, the crew of Deepwater Horizon had two main operational concerns;

1. Ensuring the 'cement job' was secure (a Gain that would allow them to move on to the next step);
2. Fracturing the formation and losing returns (a potential Loss that had been an ongoing concern throughout the operation).

It was the failure of the cement job, along with the failure of a series of controls that were expected to intervene in the event of a "kick" that led to the disaster (National Commission, 2011).

It is no more than speculation to wonder if perhaps the psychology of Prospect Theory could have influenced the perception of risk, decision-making and the application of attention and effort of the Deepwater Horizon crew on that day. But given the importance of being able to understand why the operators involved in that incident – or indeed in other psychologically equivalent situations – made the decisions and took the actions they did, it is a speculation that seems justified.

In the real world, there are always many sources of risk. The relative risk profile has to be continuously prioritised and managed in real-time. Choices have to be made about where to allocate effort and resources. Could the psychology of Prospect Theory influence how these decisions are made? Could what is proposed by Prospect Theory cause the leaders of an organization – whether consciously or unconsciously – to work twice

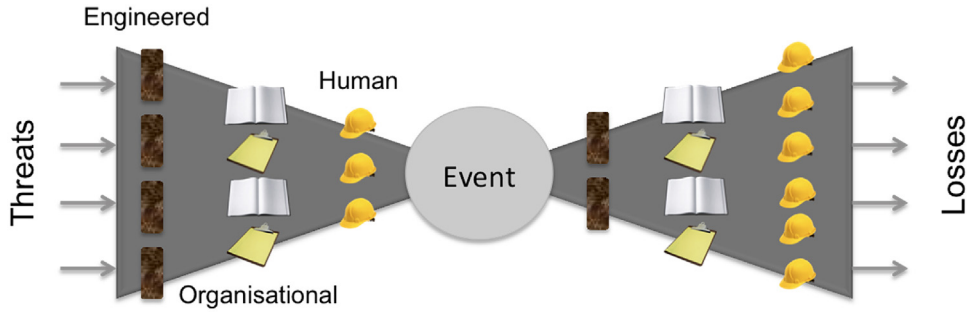

Fig. 3 – Conceptual bowtie showing the relationship between threats, events and losses and different types of generic controls (from McLeod, 2015).

as hard and focus twice as much attention and energy on the wrong things?

If Prospect Theory did indeed apply in this way, it would raise significant issues not only for ensuring the accurate prioritization of real-time risk assessments – because the largest assessed risk at any time would draw a disproportionate amount of effort and attention – but for the real-time, front-line management of operations.

3. Basic concepts in barrier management

Conceptually, the representation of layers-of-defenses is widely thought of in terms of James Reason's 'Swiss Cheese' model (Reason, 1990), where accidents occur when 'holes' in protection layers (cheese slices) align. Fig. 3, which is based on a model developed by the UK Health and Safety Executive (HSE, 2004) illustrates a related view based on a conceptual "bowtie". At the center of the bowtie (the 'knot') is an unwanted: a gas release, a fire, a dropped object, or whatever the event of concern is. The left hand side represents all of the threats that could lead to the event, while the right hand side represents the development of the event to the point where losses are incurred (injury, damage, loss of life, reputational damage, etc.). Both sides of Fig. 3 show three generic types of controls against the threats:

- The first and strongest type of control are engineered. They can involve reducing or eliminating the hazard by, for example, avoiding the use of hazardous or corrosive materials in the process. Or they can be physical barriers, such as the quality of steel, corrosion resistant paint, or mechanical or electronic interlocks.
- The second type of generic controls are organisational systems. These are the elements of the local Safety Management System, including team organization and working arrangements, job hazard assessments, procedures, work instructions, and so on put in place to control the way work is carried out.
- The third type of generic controls are human. Ensuring work is performed by trained, competent, experienced people, working in a strong safety culture, who are properly motivated and in a fit state to work.

In combination, these three types of generic controls, with potentially multiple instances of each type, provide 'layers-of-defences' against threats.

3.1. *Bowtie analysis*

There are now a number of more-or-less formalized approaches to developing and assessing the layers-of-defences assets rely on for safety and integrity. Among the most formalized and rigorous of them is probably the 'Layers of Protection Analysis' (LOPA), technique. A related technique to LOPA that is in widespread, and growing use across safety-critical industries is the technique of Bowtie analysis. The Centre for Process Safety (CCPS, 2017) is preparing guidance on conducting and using Bowtie analysis. Lewis and Smith (2010) summarise some of their experience in its application to a range of safety critical operations.

The diagrams prepared to represent the results of a Bowtie analysis usually comprise a number of elements, as illustrated on Fig. 4.

- Each diagram is associated with a specific hazard and a single top event—one of the ways in which the hazard could be released. There can be multiple top events for a single hazard.
- Threats are events that, if they are not prevented from doing so, are likely to lead to the top event occurring.
- Controls are the defences against the threat: on the left hand side of the bow-tie, they are all of those things that are considered sufficient to reduce the likelihood of the threat line leading to the top event to an acceptable level. On the right hand side, they are all of the things intended to prevent a top event, if it did occur, from leading to the consequences. (Sometimes controls on the left hand side are referred to as 'control measures', while those on the right hand side are referred to as 'recovery measures').
- Degradation factors are things that could cause a control to fail to do its intended job.
- Degradation factor controls are things that are intended to prevent the degradation factors from interfering with the functioning of the control.

When they are done properly, both LOPA and Bow-Tie Analysis provide detailed engineering analyses that identify all of the controls expected to be in place to manage the risk associated with a specific hazard. In combination, the controls included in a LOPA or Bow-Tie Analysis are expected to be sufficient to reduce the risk to a level that the organization – with, in some countries, influence from a regulator – is prepared to accept: i.e. to reduce the risk associated with a hazard to a level that is considered to be 'As Low As Reasonably Practicable' (ALARP), where the cost and effort needed to reduce the risk further is considered grossly disproportionate to the reduction in risk that would be achieved.

It seems reasonable to expect an organization that relies on a layers-of-protection strategy to be rigorous in assuring the controls they choose to rely on will actually be capable of doing the job expected of them. That applies as much to those measures that rely on human performance as to any other measures. In the UK, this is reflected in Regulation 4 of the UK Health and Safety Executive's guidance to the Control of Major

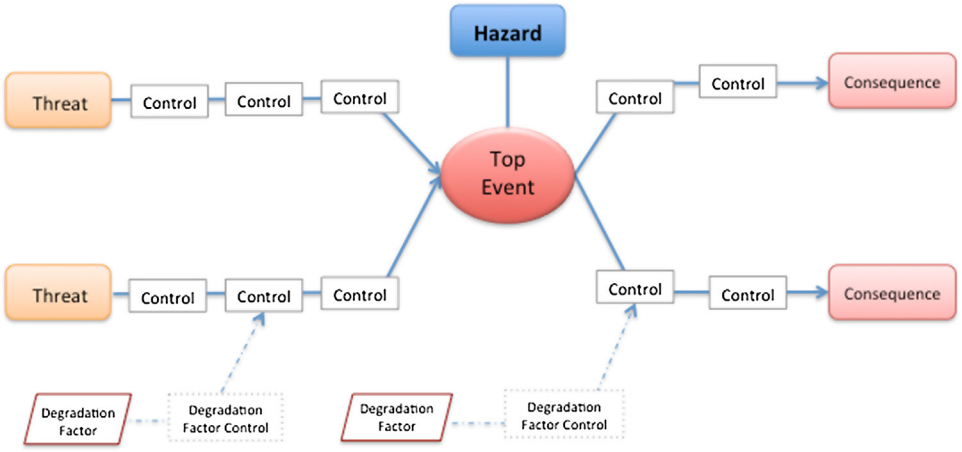

Fig. 4 – Elements of a Bowtie analysis.

Accident Hazard regulations, 1999, when it states that: '*Where reliance is placed on people as part of the necessary measures, human factor issues (including human reliability) should be addressed with the same rigour as technical and engineering measures*' (HSE, 2006, *Schedule 2, para 2*).

3.2. *Terminology*

There is inconsistency among different users of barrier models, as well as in the published literature, in the terminology used to describe the various components of barrier management systems, the nature and the relationships between them, as well as in the criteria that should be used to assess the quality of controls. Sklet (2006) provides a comprehensive review of the various forms of terminology that have been associated with the concept of barrier management systems. Hollnagel (2008) also reviews and discusses the characteristics of different approaches to barrier systems. For the purpose of this paper, Fig. 5 (from CIEHF, 2016) summarises the conceptual relationships between the various components of a barrier management system. The following definitions are used;

- The term "control" means measures expected to be in place to prevent incidents. Controls comprise barriers and safeguards.
- The term "barriers" means controls that are assessed as being sufficiently robust and reliable that they can be relied on as primary control measures against incidents. Individual barriers can be either Active or Passive, and can comprise a combination of technological and/or human elements to deliver the required barrier functionality.
- The term "safeguards" means controls that support and underpin the availability and performance of barriers but that cannot meet the standards of robustness or reliability to be relied on as a full barrier.

The distinction shown on Fig. 5 between two forms of human barriers – operational and organisational – is based on the usage of the Norwegian Petroleum Safety Agency (PSA, 2013);

- Organisational barriers are where the organisation explicitly prescribes how decisions are to be taken, and/or what is to be done by means of written rules, instructions or procedures. Decisions and actions are taken by individual operators following the prescribed instructions. There is intended to be little room for autonomy or discretion in what is done.
- Operational barriers are those where there is no specifically prescribed manner of deciding or acting. Responsibility is left to individuals having the necessary competence to take appropriate action at the time consistent with the culture, guidance, principles and constraints set by the organization. Operational barriers rely on individuals' skill and experience, capabilities in problem solving, decision making, and imagination, as well as team working skills including coordination and communication.

3.3. *Criteria for robust controls*

A number of criteria need to be met if any proposed control is to be given credit as a full barrier in a barrier management system. For process systems, IEC 61511 (IEC, 2010) requires not only that any Protection Layer must provide a minimum 100-fold reduction in risk with at least 90% availability, but it must meet four characteristics: (i) it must be specific to a single potentially hazardous event (Specificity); (ii) it must be independent of other protection layers (Independence); (iii) it can be counted on to do what it was designed to do (Dependability); and, (iv) it is capable of being audited (Auditability).

In guidance to UK industry following the fire and explosion at the Buncefield fuel storage site in 2005, the UK's Process Safety Leadership Group (PSLG, 2009) requires only three criteria: that a valid protection layer needs to be Independent, Effective and Auditable. Other organisations (CCPS, 2017; CIEHF, 2016) have recommended generally similar criteria for barriers in Bowtie analysis: though the term "Effectiveness" is generally used in place of "Dependability". Though note that for Independent Protection Layers (IPLs) included in a Layers-of-Protection-Analysis (LOPA), the Center for Chemical Process Safety (CCPS, 2015), recommends seven "core attributes" that any IPL should have: Independence; Functionality; Integrity; Reliability; Auditability; Access Security and Management of Change.

These criteria for barrier quality need careful application when they are applied to controls that rely on human performance. In particular, the human and organizational factors associated with requirements that barriers be independent and effective can be especially problematic. Issues associated with these two requirements are briefly considered in the following sub-sections.

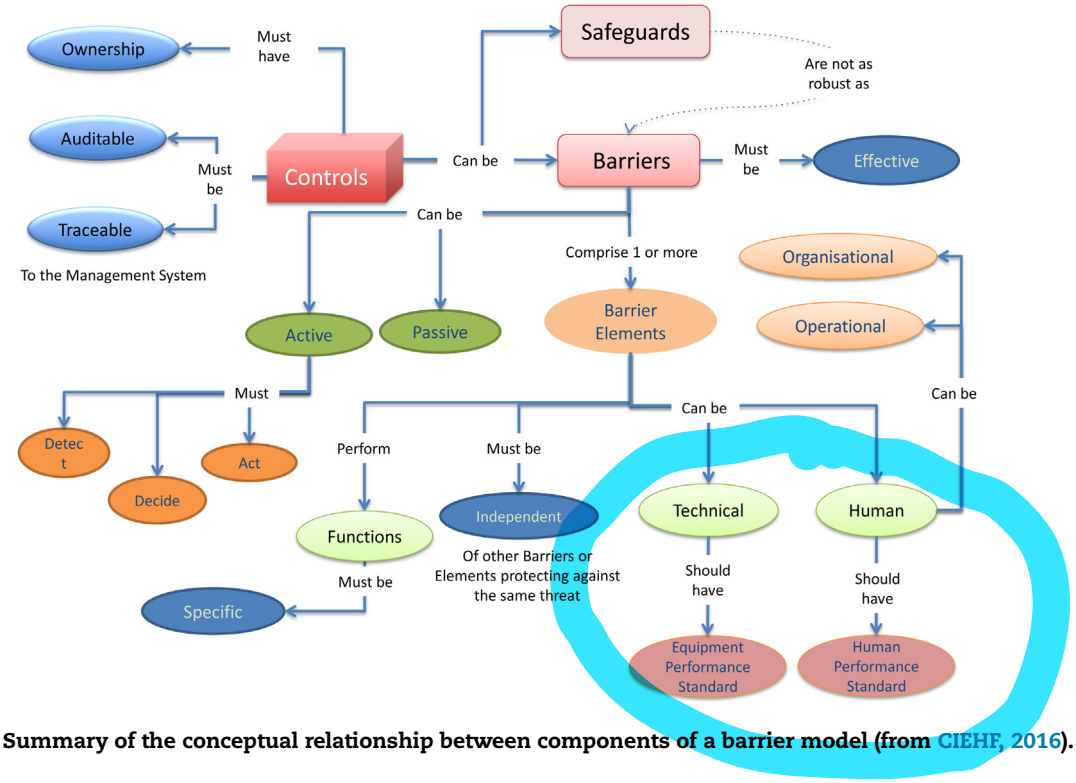

Fig. 5 – Summary of the conceptual relationship between components of a barrier model (from CIEHF, 2016).

3.3.1. *Human factors and independence*

Achieving true independence of controls in terms of their reliance on human performance can be a major challenge. Controls frequently rely on the same individual or team, so any factor – workload, fatigue, distraction, lack of training, etc. – that defeats one control will often have the potential to defeat other controls in the same threat line. And even when controls are assigned to different individuals, organisational factors – leadership messages emphasizing production over safety, rewarding individuals for delivery even when corners have been cut, poorly thought out incentive schemes that reward unsafe behavior in annual bonus schemes or the way contractors are incentivized in their contract arrangements – can all lead to the defeat of many controls that depend on people. Further, the independence that can be achieved by relying on other people to cross-check somebody's work may not be as effective as is widely assumed. This was recognised as far back as 1983 in the classic work by Swain and Guttman that has provided the basis for most approaches to quantifying human reliability since "*...the behaviour of an operator and a checker are not independent*" (Swain and Guttman, 1983). Similar learning has also been reiterated by the UK Process Safety Leadership Group: "*... the risk reduction due to checking is frequently not as great as might be expected...the intended independence of the checking process may not in fact be achieved.*" (PSLG, 2009). Unfortunately, those and similar warnings are frequently over-looked when decisions are made about how to assure human performance in operations.

3.3.2. *Human factors and effectiveness*

Making judgements about whether any control is likely to be Effective (i.e. that each control, on its own, should be capable of preventing an event from leading to an undesirable consequence) needs a great deal more information than is typically produced when a Process Hazard Analysis, HAZOP, Bowtie analysis, LOPA or other form of risk analysis to support a layers-of-defenses strategy is carried out. McLeod (2015) argues that judgements about the likely effectiveness of controls that rely on human performance means being clear about exactly what is *intended*, and what is *expected* of human performance for the control to be considered to meet the effectiveness criteria.

Intentions are defined as things that can reasonably be expected to be within the scope of influence of the team that assesses the risks and develops and approves a set of controls to mitigate those risks. Intentions will often relate to aspects of the design of the work environment or equipment interfaces. A control that relies on someone opening or closing a valve brings with it the clear intention not only that the individual will be act on the right valve (or even the right aircraft, see AAIB, 2015), but that the human interface to the valve will be designed and labelled such that the chances of design-induced human error are reduced to ALARP. Similarly, relying on an alarm as a control element brings with it numerous implicit intentions about the design quality of that alarm and the environment in which it will be presented. Expectations, by contrast, can be defined as things the analysis team cannot usually be expected to be responsible for, but that it needs to assume will be true for a control that relies on human performance to be effective.

McLeod (2015) has discussed in some depth what the requirements of Independence and Effectiveness (as well as Assurance) mean when they are applied to controls that rely on human performance. Based on application of the principles of Human Factors Engineering, and drawing on a detailed analysis of the investigation into the Buncefield explosion in 2005, McLeod explores the Human Factors challenges involved in meeting each of these requirements and demonstrates how each of them can be, and in the case of the Buncefield incident in fact were, defeated by real-world events.

3.4. *Barriers and safeguards*

Many human and organisational factors need to be managed to deliver the consistently high levels of human reliability that are expected and needed to ensure incident-free operations:

culture, organisational design, local site organisation, competence and the fitness to work of the workforce, the design of work systems and equipment interfaces, as well as contractual relationships with customers, contractors and suppliers. Weaknesses in any of these areas can lead both to failing to achieve the levels of human reliability that are expected and needed as well as increasing the chances that human performance will lead to a weakening or complete failure of barriers.

As important as these human and organisational measures are, most of them could not hope to meet the criteria to be considered as barriers in a formal barrier management system. Nevertheless they play a critical role in mitigating and managing risk, and the role they play needs to be capable of being recognized within a barrier management system.

Both the Centre for Chemical Process Safety (CCPS, 2017) and the Chartered Institute of Ergonomics and Human Factors (CIEHF, 2016) recommend that most organisational measures be treated as safeguards, rather than as barriers. Human and organizational safeguards are defined as any form of action that an organization takes, or control that it seeks to have in place, with the aim of in some way influencing behavior and reducing the chances of human error. Safeguards can range from local warnings and signs, the design and implementation of alarms and the human machine-interface to control systems, through Job Design, operating procedures and cross-checking practices, to the willingness of front-line personnel to stop work if they have any concerns over safety. Generally, the role of these organisational safeguards is to ensure that the barriers that are expected to be in place are not degraded or defeated by other factors, including human error. Safeguards cannot, and do not need, to provide the same level of risk reduction as barriers. Their role should however be recognized in any comprehensive approach to barrier analysis. Though note, as Fig. 5 indicates, that safeguards should still have clear ownership, be capable of being audited, and be traceable to some elements of the organisations management system.

In some cases – and perhaps especially in the case of human and organisational factors – whether a control is considered to be a barrier or a safeguard will come down to the willingness and ability of the organisation responsible to implement, manage and maintain barriers for specific hazardous situations. And to ensure the selected barriers are and remain, specific, independent and effective throughout the expected operational lifetime. A control which is a barrier in one situation could be treated as a safeguard elsewhere if the organisation was not able or willing to invest the necessary effort to ensure it achieved the standard necessary to be treated as a barrier in both situations.

For example, many companies invest significant time and effort trying to achieve a culture where any member of the workforce who is concerned about safety is encouraged and expected to intervene to stop the work without blame or penalty. Fig. 6(b) shows the role of "Stop culture" as a safeguard in the case of a crane lift operation. "Stop culture" could however be treated as a full barrier in other situations. The difference would be down to the willingness of the organisation involved to put the effort and resources needed into ensuring both that their STOP culture is sufficiently comprehensive, robust and justly implemented and that the relevant members of the workforce would actually intervene and stop work if they had concerns over safety. As well, of course, as the organisation's willingness to accept the increased level of disruption to operations that would inevitably go along with an increased frequency of work being stopped in the presence of what will often be weak signals of danger, many of which will be false alarms.

4. Issues with current practice

In its' White Paper on Human Factors in Barrier Management, the Chartered Institute of Ergonomics and Human Factors (CIEHF, 2016) identified eight specific concerns with the way Human and Organisational Factors are addressed in many current uses of barrier management, and Bowtie analysis in particular. In summary;

1. Top Events are frequently located too far to the right: that is, the events that barrier systems seek to avoid are too close in time to the consequences (fatalities, losses, etc.) that those events can lead to.
2. Too many "barriers" are identified, most of which are not able to meet generally accepted criteria for full barriers.
3. Barrier models rarely take a systems view of the human and organizational factors associated with the threats they are trying to control.
4. There is a lack of understanding of the nature and complexity of the tasks – and especially the cognitive elements of those tasks – that need to be carried out for barriers to function as intended.
5. There is a lack of awareness of the difference between "work-as-imagined" and "work-as-done" (Hollnagel, 2014).
6. Human "error" is commonly modelled as a threat, and proposed barriers are put in place that try to block the error from leading to a top event.
7. Intentions and expectations of human performance that are implicit in the decision to rely on people as part of a barrier system are rarely made explicit or communicated to those that need to implement, perform, support or maintain barriers.
8. Barrier models are often prepared, implemented and distributed to the workforce in a manner that does not properly support their operational use.

4.1. *Treating human error as a "threat"*

Both the CCPS (CCPS, 2017) and CIEHF (CIEHF, 2016) recognise that identifying "human error" as a threat in a Bowtie analysis can create a misleading impression that the risk of human error is being adequately managed by barriers. Problems include:

- Focusing attention on trying to minimize the risk of human error rather than recognizing the real barriers and ensuring they are as robust as they can be.
- Taking the potential for error out of the situation or context that can induce it.
- Errors identified as threats being limited to those within the experience or imagination of those involved in developing the barrier model.

Treating humans primarily as a threat also misses the opportunity to develop a deeper understanding of the ways people provide flexibility and adaptability and therefore contribute to system resilience. And it sends a fundamentally negative view of the role of people in safety management that can adversely impact development of a strong safety culture.

The alternative to treating human error as a threat is to recognize the real impact of human error: which is to defeat or degrade other barriers. Rather than focusing on the human error, attention should be directed towards improving the

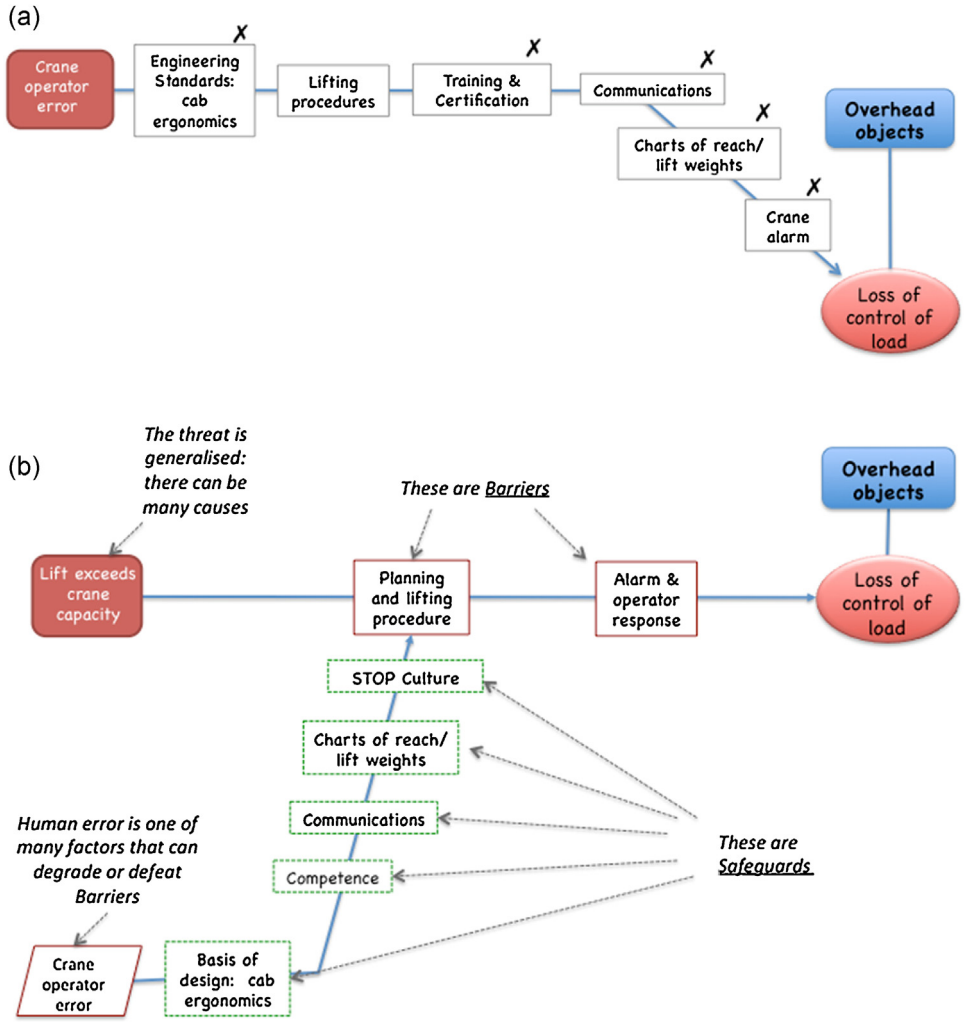

Fig. 6 – (a) Left-hand side of a Bowtie analysis for loss of control of a load during a crane lift showing human error as a threat. Most of the barriers do not satisfy the effectiveness criteria for barriers. (b) Alternate bowtie for loss of control during a crane lift where human error is treated as one of many factors that can degrade barriers. There are only two genuine barriers, but many potential safeguards to prevent human error from degrading those barriers.

inherent strength and resilience of the barrier(s) that the error could defeat or degrade.

To illustrate this alternative approach, Fig. 6(a) shows a representation of the left-hand side of an issued bowtie diagram for the hazard posed by overhead objects during heavy-lift crane operations. In this case, the company had identified "crane operator error" as a primary threat that could lead to loss of control of the load during the lift. The specific error addressed was that the crane operator attempted to lift a load that exceeded the lifting ability of the crane. The company proposed the six controls shown on the diagram as barriers to block the threat path. Even a cursory examination of these barriers quickly makes clear that five of them (shown with an x on the figure) could not, on their own, meet the effectiveness criteria. (Whether "Lifting procedure" could meet the effectiveness criteria is also debatable, depending on how the requirement for active barriers to have detect, decide and act functionality is interpreted).

Fig. 6(b) shows an alternative treatment of this same hazardous situation. In this case, the threat has been generalised away from human error, to cover any situation where the load exceeds the crane's capacity. Rather than simply focusing on error on the part of the crane operator, this would cover, for example, situations involving errors in labeling of the item to be lifted, or communications either between the supplier and the organisation carrying out the lift or within the local workforce. Fig. 6(b) shows two active barriers: (i) Planning the lift and following an approved lifting procedure, and (ii) an overweight alarm and the associated operator response. Planning could meet the functionality of detecting what needs to be done and deciding how it is to be achieved, while following an approved procedure could meet the "act" requirement. These two elements, in combination, therefore provide the functionality needed of a full active barrier. The original barrier "crane alarm" has been expanded to include the operator response that adds the "decide" and "act" functionality to the "detect" function provided by the automated alarm system.

Fig. 6(b) shows crane operator error as one (of potentially many) factors that has the potential to defeat or degrade the "planning and lifting procedure" barrier. (It could also, of course, defeat the "alarm and operator response", though via different mechanisms). Finally, the figure shows five safeguards that could be expected to contribute to reducing the likelihood of an operator error from defeating the planning and procedure barrier. Note that four of these safeguards were originally considered, on Fig. 6(a), to be barriers. They have an important role to play in protecting this hazardous operation from risk. But they cannot meet the standards needed to be considered barriers, and are therefore treated as safeguards:

they are important, but they do not provide the same level of protection as the full barriers.

Fig. 6(b) also recognises that the existence of a "STOP culture" is a safeguard. As noted earlier, many companies try to achieve a culture where any member of the workforce who is concerned about safety is encouraged and expected to intervene to stop the work without blame or penalty. Recognising the role of a STOP culture as a safeguard formalises this aim. It makes clear to everyone who develops, authorises, accesses or uses the bowtie – from senior management to the front-line workforce – the place that STOP culture has in the safe conduct of operations at the front-line. The same is true of many other organisational controls: from the use of engineering standards, contractor management and incentive schemes to shift planning, job design and many others. Representing them explicitly as safeguards on bowtie diagrams makes their intended role in safety management clear.

5. Example incidents

To demonstrate some of the insight that can come from examining the implicit intentions and expectations about human performance that underpin assumptions about control effectiveness, as well as the ways in which the "hard truths" of human performance can defeat barriers, the remainder of this paper will consider two incidents that occurred in process systems.

5.1. *Over-pressurised pig launcher*

A team was preparing for a pipeline inspection using an in-line inspection tool (known as a 'pig'). The team believed that the pipeline valves were open and began pumping nitrogen from a truck to purge the line. However, the valve between the pig-launcher and the pipeline was actually closed preventing nitrogen from entering the pipeline. The truck included a pressure trip set at 6000 psi although the Maximum Allowable Working Pressure (MAWP) of the pig launcher was 350 psi (i.e. the truck was capable of supplying nitrogen at a much high pressure than the launcher was designed to withstand). The pig launcher was not equipped with a pressure relief valve. When pressure was applied, the 100-psi gauge on the pig launcher almost instantaneously swung to the zero position. The team at the pig launcher mistakenly read the gauge as indicating there was no nitrogen flowing from the nitrogen truck and called for more pressure. The pressure release happened within two minutes of the call from increased nitrogen flow. One operator was killed and two others were hospitalised.

What kind of controls might have been in place that should have prevented the pig launcher from being over-pressurised? The obvious ones would include a pressure trip and relief valves. More importantly, what sort of expectations might the organisation involved have reasonably held about those controls? Here are two suggestions that can reasonably be assumed from the incident report:

- The job would have been planned and a safety review conducted before starting work;
- Operators would be aware of the risks and exercise caution. If they had any doubt they would stop.

It must have been assumed that holding a safety review prior to starting work would identify the risks, and ensure everyone involved knew what was involved in ensuring they were controlled. Safety reviews prior to starting work are widely relied across many industries. A safety review had indeed been carried out (the control was in place), but it was not effective in recognising or mitigating the over-pressurisation risk. So why did it fail in this case? Was there something unusual about the way the review was held on that day, or the engagement of the people involved at the time? Are safety reviews usually effective as a means of raising awareness and ensuring risks are under control? Or was this only one incident among many where such safety reviews have failed to identify and provide the expected control over significant risks? Are they, in fact, as effective as is widely assumed?

Or consider the expectation that the operators would be aware of the risks associated with pig launchers, will exercise caution, and will stop the job if they are in any doubt. The evidence suggested that this expectation was not met. The individual(s) who read the pressure gauge, concluded that there was no pressure in the pig launcher and called for the flow of nitrogen to be increased cannot have been in any doubt. They did not mis-believe the pressure gauge: they believed it was reading zero flow, when in reality the pressure was already too high for it to provide a reading. But there can have been no doubt involved, and therefore no need to exercise caution and stop and question what they were about to do. It has the characteristics of decision making dominated by System 1 thinking.

5.2. *A Pipeline corrosion incident*

Organisations have to believe that at the time a risk analysis is performed their analysis teams are capable of identifying the ways operations can be put at risk as a consequence of human performance – which can be many years before an event actually takes place. Because of the situational nature of human performance, and the broad range of factors that influence how any individual perceives, understands and experiences the situation they are in at the time they assess a situation, make decisions and act, that belief can be extremely optimistic. The following incident illustrates how some of the hard truths of human performance discussed earlier can play out to defeat unrealistic expectations about how people will behave and perform in the real world.

A fire occurred when a sample supply line was breached due to corrosion in a carbon steel line allowing the release of hydrocarbons to atmosphere under high pressure and at high temperature. A change in sampling procedures had been introduced some time previously that required the valves on the sample line to be accessed more frequently than had been anticipated during design. The valves were difficult to access, being located up to 15 feet above the standing position, and were known to be difficult to operate. With the increased sampling frequency, the practice became to leave the valves open, avoiding difficulties operating the valves. As a result, the piping became more exposed to the corrosive environment in the process stream, accelerating the rate of corrosion and leading to the breach.

At the heart of this incident was an operational change that increased the frequency of operating a valve that was both inaccessible and difficult to work with. That change led operators to adopt a working practice that defeated the premises of the barrier strategy. Allow, for the purpose of this discussion, that access to the valves had been optimized during design in accordance with a Valve Criticality Analysis (see ASTM, 2013 for a description of the Valve Criticality Analysis

method) reflecting the anticipated frequency of use and criticality of the valves. As was discussed earlier, one of the hard truths of human performance is that if working life is made unnecessarily difficult, people will find an easier way to do things (even if it is more risky). Would a PHA, HAZOP or LOPA analysis be capable of recognizing that a change in sampling policy could trigger that hard truth, leading operators to adopt a working practice that defeated key assumptions behind the controls? Would a change in sampling policy be subject to a formal Management of Change procedure? One that not only recognized the implications for operators working with inaccessible and difficult to use valves, but that considered the possible implications on human behavior if the change made work more difficult? As Myers has pointed out: "*. . .unintended or poorly managed changes may be a more important consideration than those due to malicious intent*", (Myers, 2013).

The piping design would certainly have allowed for the anticipated rates of corrosion based on the anticipated exposure to the process materials. And the corrosion inspection plan would have been based on the anticipated exposure rate, assuming that the valves would only be open during sampling: a significantly lower exposure to the process stream than had been introduced.

Even for a relatively simple incident such as this, there are major challenges facing any team tasked, *a priori*, with identifying the human performance issues that could lead to an incident. And it is clear that a simple bowtie or LOPA analysis cannot capture what the organization really would have intended and expected should have prevented this incident. While it is easy, with hindsight, to see how the holes in the Swiss cheese lined up in this situation, that is very far from the case when trying to do an analysis proactively, and to identify the likely situations that could realistically give rise to a threat scenario. Indeed, it might be thought unrealistic to expect any analysis team to be capable of coming up with a proactive analysis that actually captures what did in reality happen in this incident.

The question is whether it is reasonable to expect any risk analysis team to imagine how the combination of events that actually did occur in this incident might coincide. And, even if they had the imagination, is it likely that they would give it credibility as being a realistic combination of events that needed controls in place? It is far more likely that operators would simply be assumed to be competent, and expected to follow procedures.

6. The human performance specification

It is normal practice for barriers that are purely technical in nature to have an Equipment Performance Standard associated with them that specifies the characteristics and performance required. According to the Norwegian Petroleum Safety Agency (PSA): "*The most important consideration is not the label attached to the various barrier elements, but the presence of identified and established performance requirements for all the barrier elements regarded as necessary to implement the barrier functions*" (PSA, 2013). A performance standard specifies the objective, measureable performance and assurance or verification steps required for that barrier. As Fig. 5 shows, the same concept applies equally to barriers that are, or that include, human barrier elements.

A Human Performance Standard for barriers, or barrier elements should cover at least seven topics (CIEHF, 2016):

1. In what way is the performance the barrier will deliver specific to the threat and the situation when the barrier function is needed;
2. Who is involved in delivering the required performance? That includes:
   - Who detects that the barrier function is needed?
   - Who decides what is to be done?
   - Who takes action?
3. What information is needed for successful performance of the function?
4. What decisions or judgements are likely to be involved?
5. What actions need to be taken, and how will the operators involved know whether the actions have been completed successfully?
6. Any Human Factors Engineering or other technical standards that are expected to be complied with in the design and layout of work systems, equipment interfaces or support systems (including procedures and work instructions) associated with performance of each of the detect, decide or act elements of the barrier function.
7. The standard for successful performance of the barrier. Performance criteria could, for example, cover;
   - The maximum allowable time to detect an event expected to trigger the function
   - Accuracy of interpreting the event
   - The maximum allowable time to initiate an intervention
   - The maximum allowable time to complete the intervention
   - Maximum acceptable number or percentage of missed events (i.e. failing to perform the barrier function when it should have been performed)
   - Maximum acceptable number or percentage of false alarms (i.e. performing the barrier function when it was not needed)
   - Tolerance limits for acceptable performance

A Human Performance Standard should also document any specific expectations about how operations around the barrier will be conducted that are especially critical to performing its function. Table 1 illustrates how a Human Performance Standards for the barrier "Planning and lifting procedure" in the bowtie for loss of control of a load during a crane operation (Fig. 6(b)) could be documented.

7. Conclusions

An issued layers of defenses strategy, in whatever form it takes, is a significant statement of intent on behalf of the organisation that prepares and approves it. Organisations can make improvement in assuring the human controls they rely on for safety and production are as robust as they reasonably can be by being clearer about exactly what is the role of people in their systems. When controls that rely on human performance are relied on, they should be subject to robust challenge about what exactly is intended and what is expected for those controls to do what is needed of them.

Some people in positions of responsibility find it difficult to treat the "hard-truths" of human behavior and performance seriously in the decisions they make and the actions they take. They believe people who are trained and competent, are assessed as being fit for work, properly motivated and working in a culture that places a high value on safety and compliance, should in some way be capable of overcoming not only poorly designed working environments and procedures, but very powerful instincts of human nature. Though the reality

Table 1 – Example human performance standard for human barrier "planning and lifting procedure" (see Fig. 5) (from CIEHF, 2016).

Barrier	Planning and lifting procedure
Barrier element	Lifting procedure Type Organisational
Barrier function(s)	Plan, prepare and carry out crane lifts in accordance with company standard xyz
Limits	The barrier is intended to provide protection for lifts carried out from a fixed base using mobile cranes with loads in the range from X to Y tonnes.
Active or Passive?	Active
What makes the barrier specific to the threat?	Lifting procedure ABC is specific to mobile cranes of type NNN, and to lifts performed from a fixed base with loads in the range from X to Y tonnes.
Performance Criteria	Drivers should be able to; 1. Access lifting procedure ABC without having to leave their cab. 2. Comply with all of the steps in the procedure, using only charts of reach/lift weights if necessary. 3. Recognise when any proposed lift is outside the scope of the procedure.
Timing	From the point where a load is clear of the ground, all lifts should be able to be completed without any change of crew, and before any significant change in weather conditions.
Who is involved?	Lifting supervisor; Crane driver; Banksmen
Who Detects?	Supervisor and crane driver should know when the Lifting Procedure is to be followed.
Who Decides?	Decisions on lifts should be taken by the Crane Driver in compliance with the lifting procedure and charts of reach/lift weights
Who Acts?	Crane driver, supported by Lifting Supervisor and Banskmen.
Information needed	• Crane capability • Crane location • Location of lift and lay-down areas • Details of lifts • Lift route • Weather forecast for the lift period • Availability and experience of banksmen
Key judgements or decisions involved?	• Whether any lifts are likely to approach safe lift limits • Reliability of weather forecast for duration of lift • Whether required lifts are within driver experience and competence • Whether there is sufficient manpower available
Actions	• Carry out lift in accordance with lifting procedure
Feedback	Feedback to the crane driver of the state of the lift achieved by; a) direct visual monitoring of the load, b) visual monitoring of in-cab instruments, c) audio monitoring of radio communications between the crew, d) visual monitoring of banksmen's hand-signals.
Engineering standards	• Cab ergonomics, including visibility and viewing angles to be compliance with ISO XXX • Lifting accessories (hooks, shackles, link chains, etc.) to be compliant with XXX • Lifting points designed onto major items to be compliant with XXX.
Critical Expectations associated with human performance for the barrier to be effective.	
Company lifting standard xyz will be up to-date and a current version available in the crane cab.	
The lifting standard will have been subject to a Procedural HAZOP to ensure it is fit for use in a safety-critical role.	
Contractors will have no commercial or personal incentives not to comply with the plan.	
The driver and crew will understand the importance of complying with the lifting procedure and will advise line management if they have any concerns either about its suitability or with the way it is being implemented.	
Operational crew will not go ahead with lifts if the conditions of the lift (such as crane type, supporting structure or nature of the loads) are outside the limits of the standard.	

is that we are all human, and are all subject to these powerful motivations.

As the Nobel prize winning psychologist Daniel Kahneman puts it, "*. . .people will eventually gravitate to the least demanding course of action. . .laziness is built deep into our nature*" (Kahneman, 2012). That is not to suggest, when incidents do occur, that people should not be held responsible for violating what are clear expectations and procedures. But it does mean that, when layers-of-defenses strategies are being developed, they should be subject to robust challenge about whether the expectations of human behavior and performance that are relied on for those barriers to achieve an acceptable level of effectiveness and reliability are realistic and reasonable. And when changes to operational practices are being implemented, the potential consequences on human behaviour need to be considered carefully.

People are nearly always a positive element in complex socio-technical systems. The objective should therefore be to strive to make people as reliable as possible. Organisations operating complex socio-technical systems should seek to ensure they have in place the necessary systems and support structures, and should design and operate their activities in ways that allows people to be as productive and adaptable as they can be. That can mean shifting from a mindset that focuses on ensuring the risk of human error is "As Low As Reasonably Practicable" – ALARP – and towards one of ensuring operations and work systems are designed and operated in such a way that human reliability can be "As High As Reasonably Practical" – AHARP.

References

AAIB, 2015. Report on the Accident to Airbus A39-131, G-EUOE London Heathrow Airport 24 May 2013. Aircraft Accident Report 1/2015. Air Accidents Investigation Branch.

ASTM, 2013. Standard Practice for Human Engineering Design for Marine Systems, Equipment and Facilities. ASTM International.

CCPS, 2015. Guidelines for Initiating Events and Independent Protection Layers in Layer of Protection Analysis. Center for Chemical Process Safety, Wiley & sons, New Jersey.

CCPS, 2017. Guidelines for Bowtie Risk Management. Center for Chemical Process Safety, in press.

Chambers, C., Willday, J., Turner, S., 2009. A Review of Layers-of-protection Analysis (LOPA) of Overfill of Fuel Storage Sites. HSE Books.

Chemical Safety Board, 2016. Investigation Report Volume 3: Drilling Rig Explosion and Fire at the Macondo Well. Report no. 2010-10-I-OS. Chemical Safety Board.

CIEHF, 2016. Human Factors in Barrier Management. Chartered Institute of Ergonomics and Human Factors.

Hamilton, I.W., Turner, C., 2014. Building a Culture of Effective Process Safety Management SPE-172323-MS. Society of petroleum Engineers.

HSE, 2004. Human Factors Guidance for Selecting Appropriate Maintenance Strategies for Safety in the Offshore Oil and Gas Industry. Research Report 213. Health and Safety Executive.

HSE, 2006. A Guide to the Control of Major Accident Hazards Regulations 1999 (as amended): Guidance on Regulations. HSE Books.

Hollnagel, E., 2008. Risk + barriers = safety? Saf. Sci. 46, 221–229.

Hollnagel, E., 2014. Safety-I and Safety-II: The Past and Future of Safety Management. Ashgate, Farnham.

IEC, 2003. Functional Safety of Electrical/Electronic/Programmable Electronic Safety-related Systems. IEC 61508. International Electrotechnical Commission.

IEC, 2010. Functional Safety—Safety Instrumented Systems for the Process Industry Sector. IEC 61511. International Electrotechnical Commission.

Kahneman, D., 2012. Thinking, Fast and Slow. Allen Lane, London.

Lewis, S., Smith, K., 2010. Lessons learned from real world application of the bow-tie method. In: American Institute of Chemical Engineers 6th Global Congress on Process Safety, San Antonio, Texas.

McLeod, R.W., 2015. Designing for Human Reliability: Human Factors Engineering for the Oil, Gas and Process Industries. Gulf Professional Publishing, Oxford.

McLeod, R.W., 2016a. Issues in assuring human controls in layers of defences strategies. Chem. Eng. Trans. 48, 925–930, http://dx.doi.org/10.3303/CET1648155.

McLeod, R.W., 2016b. The Impact of Styles of Thinking and Cognitive Bias on How People Assess Risk and Make Real World Decisions. SPE Paper 179197-PA. Society of Petroleum Engineers.

McLeod, R.W., 2016c. Implications of Styles of Thinking for Risk Awareness and Decision-making in Safety Critical Operations. Cognitia, September. Human Factors and Ergonomics Society.

MMS, 2005. Human Engineering Factors Result in Increasing Number of Riser Disconnects. Safety Alert No. 231. US Department of the Interior Minerals Management Service, Gulf of Mexico OCS Region.

Myers, P.M., 2013. Layer of protection analysis—quantifying human performance in initiating events and independent protection layers. J. Loss Prev. Process Ind. 26, 534–546.

National Commission on the BP Deepwater Horizon Oil Spill and Offshore Drilling, 2011. Deepwater: The Gulf Oil Disaster and the Future of Offshore Drilling: Report to the President. National Commission.

PSA, 2013. Principles for barrier management in the petroleum industry. Petroleum Safety Agency http://www.psa.no/getfile.php/PDF/Barrierenotatet%202013%20engelsk%20april.pdf.

PSLG, 2009. Safety and Environmental Standards for Fuel Storage Sites. Process Safety Leadership Group, HSE Books.

RAIB, 2015. Fatal Accident Involving a Track Worker Near Newark North Gate Station, 22 January 2014. Report 01/2015. Rail Accident Investigation Branch.

Reason, J., 1990. Human error. Cambridge University Press, New York.

Sklet, S., 2006. Safety barriers: definition, classification, and performance. J. Loss Prev. Process Ind. 19, 494–506.

Swain, A.D., Guttman, H.E., 1983. Handbook of Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications. Final Report. NUREG/CR-1278. USNRC.

Tversky, A., Kahneman, D., 1974. Judgement under uncertainty: heuristics and biases. Science 185, 1124–1131.